Analysis of Variance

Table Of Contents

Analysis of Variance

Overview

Analysis of Variance Overview

Analysis of variance (ANOVA) is similar to regression in that it is used to investigate and model the relationship between a response variable and one or more independent variables. However, analysis of variance differs from regression in two ways: the independent variables are qualitative (categorical), and no assumption is made about the nature of the relationship (that is, the model does not include coefficients for variables). In effect, analysis of variance extends the two-sample t-test for testing the equality of two population means to a more general null hypothesis of comparing the equality of more than two means, versus them not all being equal. Several of Minitab's ANOVA procedures, however, allow models with both qualitative and quantitative variables.

Minitab's ANOVA capabilities include procedures for fitting ANOVA models to data collected from a number of different designs, for fitting MANOVA models to designs with multiple response, for fitting ANOM (analysis of means) models, and graphs for testing equal variances, for confidence interval plots, and graphs of main effects and interactions.

ANOVA

Stat > ANOVA

Allows you to perform analysis of variance, test for equality of variances, and generate various plots.

Select one of the following commands:

One-Way – performs a one-way analysis of variance, with the response in one column, subscripts in another and performs multiple comparisons of means

One-Way (Unstacked) – performs a one-way analysis of variance, with each group in a separate column

Two-way – performs a two-way analysis of variance for balanced data

Analysis of Means – displays an Analysis of Means chart for normal, binomial, or Poisson data

Balanced ANOVA – analyzes balanced ANOVA models with crossed or nested and fixed or random factors

General Linear Model – analyzes balanced or unbalanced ANOVA models with crossed or nested and fixed or random factors. You can include covariates and perform multiple comparisons of means.

Fully Nested ANOVA – analyzes fully nested ANOVA models and estimates variance components

Balanced MANOVA – analyzes balanced MANOVA models with crossed or nested and fixed or random factors

General MANOVA – analyzes balanced or unbalanced MANOVA models with crossed or nested and fixed or random factors. You can also include covariates.

Test for Equal Variances – performs Bartlett's and Levene's tests for equality of variances

Interval Plot – produces graphs that show the variation of group means by plotting standard error bars or confidence intervals

Main Effects Plot – generates a plot of response main effects

Interactions Plot – generates an interaction plots (or matrix of plots)

More complex ANOVA models

Minitab offers a choice of three procedures for fitting models based upon designs more complicated than one- or two-way designs. Balanced ANOVA and General Linear Model are general procedures for fitting ANOVA models that are discussed more completely in Overview of Balanced ANOVA and GLM.

- Balanced ANOVA performs univariate (one response) analysis of variance when you have a balanced design (though one-way designs can be unbalanced). Balanced designs are ones in which all cells have the same number of observations. Factors can be crossed or nested, fixed or random. You can also use General Linear Models to analyze balanced, as well as unbalanced, designs.
- General linear model (GLM) fits the general linear model for univariate responses. In matrix form this model is Y = XB + E, where Y is the response vector, X contains the predictors, B contains parameters to be estimated, and E represents errors assumed to be normally distributed with mean vector 0 and variance σ. Using the general linear model, you can perform a univariate analysis of variance with balanced and unbalanced designs, analysis of covariance, and regression. GLM also allows you to examine differences among means using multiple comparisons.
- Fully nested ANOVA fits a fully nested (hierarchical) analysis of variance and estimates variance components. All factors are implicitly assumed to be random.

Special analytical graphs

- Test for equal variances performs Bartlett's (or F-test if 2 levels) and Levene's hypothesis tests for testing the equality or homogeneity of variances. Many statistical procedures, including ANOVA, are based upon the assumption that samples from different populations have the same variance.
- Interval plot creates a plot of means with either error bars or confidence intervals when you have a one-way design.
- Main effects plot creates a main effects plot for either raw response data or fitted values from a model-fitting procedure. The points in the plot are the means at the various levels of each factor with a reference line drawn at the grand mean of the response data. Use the main effects plot to compare magnitudes of marginal means.
- Interactions plot creates a single interaction plot if two factors are entered, or a matrix of interaction plots if 3 to 9 factors are entered. An interactions plot is a plot of means for each level of a factor with the level of a second factor held constant. Interactions plots are useful for judging the presence of interaction, which means that the difference in the response at two levels of one factor depends upon the level of another factor. Parallel lines in an interactions plot indicate no interaction. The greater the departure of the lines from being parallel, the higher the degree of interaction. To use an interactions plot, data must be available from all combinations of levels.

Use Factorial Plots generate main effects plots and interaction plots specifically for 2-level factorial designs, such as those generated by Create Factorial Design and Create RS Design.

Examples of ANOVA

Minitab Help offers examples of the following analysis of variance procedures:

One-way Analysis of Variance: Stacked Data

Two-way Analysis of Variance

Analysis of Means: Two-Way with Normal Data

Analysis of Means: Binomial response data

Analysis of Means: Poisson response data

ANOVA: Two Crossed Factors

ANOVA: Repeated Measures Design

ANOVA: Mixed Model with Restricted and Unrestricted Cases

GLM: Multiple comparisons with an unbalanced nested design

GLM: Fitting linear and quadratic effects

Fully Nested ANOVA

Balanced MANOVA

Test for Equal Variances

Interval Plot

Main Effects Plot

Interaction Plots: with Two Factors

Interaction Plots: with more than Two Factors

References for ANOVA

[1] R.E. Bechhofer and C.W. Dunnett (1988). "Percentage points of multivariate Student t distributions," *Selected Tables in Mathematical Studies*, Vol.11. American Mathematical Society.

[2] M.B. Brown and A.B. Forsythe (1974). "Robust Tests for the Equality of Variance,"*Journal of the American Statistical Association*, 69, 364–367.

[3] H.L. Harter (1970). *Order Statistics and Their Uses in Testing and Estimation*, Vol.1. U.S. Government Printing Office.

[4] A.J. Hayter (1984). "A proof of the conjecture that the Tukey-Kramer multiple comparisons procedure is conservative," *Annals of Statistics*, 12, 61–75.

[5] D.L. Heck (1960). "Charts of Some Upper Percentage Points of the Distribution of the Largest Characteristic Root," *The Annals of Statistics*, 625–642.

[6] C.R. Hicks (1982). *Fundamental Concepts in the Design of Experiments*, Third Edition. CBC College Publishing.

[7] Y. Hochberg and A.C. Tamhane (1987). *Multiple Comparison Procedures*. John Wiley & Sons.

[8] J.C. Hsu (1984). "Constrained Two-Sided Simultaneous Confidence Intervals for Multiple Comparisons with the Best," *Annals of Statistics*, 12, 1136–1144.

[9] J.C. Hsu (1996). *Multiple Comparisons, Theory and methods.* Chapman & Hall.

[10] R. Johnson and D. Wichern (1992). *Applied Multivariate Statistical Methods*,Third Edition. Prentice Hall.

[11] H. Levene (1960). *Contributions to Probability and Statistics*. Stanford University Press, CA.

[12] T.M. Little (1981). "Interpretation and Presentation of Result," *HortScience*, 19, 637-640.

[13] G.A. Milliken and D.E. Johnson (1984). *Analysis of Messy Data,* Volume I. Van Nostrand Reinhold.

[14] D.C. Montgomery (1991). *Design and Analysis of Experiments*, Third Edition. John Wiley & Sons.

[15] D. Morrison (1967). *Multivariate Statistical Methods.* McGraw-Hill.

[16] L.S. Nelson (1974). "Factors for the Analysis of Means," *Journal of Quality Technology*, 6, 175–181.

[17] L.S. Nelson (1983). "Exact Critical Values for Use with the Analysis of Means", *Journal of Quality Technology,* 15, 40–44.

[18] P.R. Nelson (1983). "A Comparison of Sample Sizes for the Analysis of Means and the Analysis of Variance," *Journal of Quality Technology*, 15, 33–39.

[19] J. Neter, W. Wasserman and M.H. Kutner (1985). *Applied Linear Statistical Models*, Second Edition. Irwin, Inc.

[20] R.A. Olshen (1973). "The conditional level of the F-test," *Journal of the American Statistical Association*, 68, 692–698.

[21] E.R. Ott (1983). "Analysis of Means–A Graphical Procedure," *Journal of Quality Technology*, 15, 10–18.

[22] E.R. Ott and E.G. Schilling (1990). *Process Quality Control–Troubleshooting and Interpretation of Data*, 2nd Edition. McGraw-Hill.

[23] P.R. Ramig (1983). "Applications of the Analysis of Means," *Journal of Quality Technology*, 15, 19–25.

[24] E.G. Schilling (1973). "A Systematic Approach to the Analysis of Means," *Journal of Quality Technology*, 5, 93–108, 147–159.

[25] S.R. Searle, G. Casella, and C.E. McCulloch (1992). *Variance Components.* John Wiley & Sons.

[26] N.R. Ullman (1989). "The Analysis of Means (ANOM) for Signal and Noise," *Journal of Quality Technology*, 21, 111–127.

[27] E. Uusipaikka (1985). "Exact simultaneous confidence intervals for multiple comparisons among three or four mean values," *Journal of the American Statistical Association*, 80, 196–201.

[28] B.J. Winer (1971). *Statistical Principles in Experimental Design*, Second Edition. McGraw-Hill.

Acknowledgment

We are grateful for assistance in the design and implementation of multiple comparisons from Jason C. Hsu, Department of Statistics, Ohio State University and for the guidance of James L. Rosenberger, Statistics Department, The Pennsylvania State University, in developing the Balanced ANOVA, Analysis of Covariance, and General Linear Models procedures.

One-Way

One-way and two-way ANOVA models

- One-way analysis of variance tests the equality of population means when classification is by one variable. The classification variable, or factor, usually has three or more levels (one-way ANOVA with two levels is equivalent to a t-test), where the **level** represents the treatment applied. For example, if you conduct an experiment where you measure durability of a product made by one of three methods, these methods constitute the levels. The one-way procedure also allows you to examine differences among means using multiple comparisons.
- Two-way analysis of variance performs an analysis of variance for testing the equality of populations means when classification of treatments is by two variables or factors. In two-way ANOVA, the data must be balanced (all cells must have the same number of observations) and factors must be fixed.

  If you wish to specify certain factors to be random, use Balanced ANOVA if your data are balanced; use General Linear Models if your data are unbalanced or if you wish to compare means using multiple comparisons.

One-Way Analysis of Variance

Stat > ANOVA > One-way

Performs a one-way analysis of variance, with the dependent variable in one column, subscripts in another. If each group is entered in its own column, use Stat > ANOVA > One-Way (Unstacked).

You can also perform multiple comparisons and display graphs of your data.

Dialog box items

Response: Enter the column containing the response.

Factor: Enter the column containing the factor levels.

Store Residuals: Check to store residuals in the next available column.

Store fits: Check to store the fitted values in the next available column.

Confidence level: Enter the confidence level. For example, enter 90 for 90%. The default is 95%.

<Comparisons>

<Graphs>

Data – One-Way with Stacked Data

The response variable must be numeric. Stack the response data in one column with another column of level values identifying the population (stacked case). The factor level (group) column can be numeric, text, or date/time. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. See Ordering Text Categories. You do not need to have the same number of observations in each level. You can use Make Patterned Data to enter repeated factor levels.

Note If your response data are entered in separate worksheet columns, use Stat > ANOVA > One-Way (Unstacked).

To perform a one-way analysis of variance with stacked data

1 Choose **Stat > ANOVA > One-Way**.

2 In **Response**, enter the column containing the response.

3 In **Factor**, enter the column containing the factor levels.

4 If you like, use any dialog box options, then click **OK**.

Discussion of Multiple Comparisons

The multiple comparisons are presented as a set of confidence intervals, rather than as a set of hypothesis tests. This allows you to assess the practical significance of differences among means, in addition to statistical significance. As usual, the null hypothesis of no difference between means is rejected if and only if zero is not contained in the confidence interval.

The selection of the appropriate multiple comparison method depends on the desired inference. It is inefficient to use the Tukey all-pairwise approach when Dunnett or MCB is suitable, because the Tukey confidence intervals will be wider and the hypothesis tests less powerful for a given family error rate. For the same reasons, MCB is superior to Dunnett if you want to eliminate levels that are not the best and to identify those that are best or close to the best. The choice of Tukey versus Fisher methods depends on which error rate, family or individual, you wish to specify.

Individual error rates are exact in all cases. Family error rates are exact for equal group sizes. If group sizes are unequal, the true family error rate for Tukey, Fisher, and MCB will be slightly smaller than stated, resulting in conservative confidence intervals [4,22]. The Dunnett family error rates are exact for unequal sample sizes.

The results of the one-way F-test and multiple comparisons can conflict. For example, it is possible for the F-test to reject the null hypothesis of no differences among the level means, and yet all the Tukey pairwise confidence intervals contain zero. Conversely, it is possible for the F-test to fail to reject, and yet have one or more of the Tukey pairwise confidence intervals not include zero. The F-test has been used to protect against the occurrence of false positive differences in means. However, Tukey, Dunnett, and MCB have protection against false positives built in, while Fisher only benefits from this protection when all means are equal. If the use of multiple comparisons is conditioned upon the significance of the F-test, the error rate can be higher than the error rate in the unconditioned application of multiple comparisons [15].

Comparisons – One-Way Multiple Comparisons with Stacked Data

Stat > ANOVA > One-Way > Comparisons

Provides confidence intervals for the differences between means, using four different methods: Tukey's, Fisher's, Dunnett's, and Hsu's MCB. Tukey and Fisher provide confidence intervals for all pairwise differences between level means. Dunnett provides a confidence interval for the difference between each treatment mean and a control mean. Hsu's MCB provides a confidence interval for the difference between each level mean and the best of the other level means. Tukey, Dunnett and Hsu's MCB tests use a family error rate, whereas Fisher's LSD procedure uses an individual error rate.

Which multiple comparison test to use depends on the desired inference. It is inefficient to use the Tukey all-pairwise approach when Dunnett or Hsu's MCB is suitable, because the Tukey confidence intervals will be wider and the hypothesis tests less powerful for a given family error rate. For the same reasons, Hsu's MCB is superior to Dunnett if you want to eliminate levels that are not the best and to identify those that are best or close to the best. The choice of Tukey versus Fisher depends on which error rate, family or individual, you wish to specify.

Dialog box items

Tukey's, family error rate: Check to obtain confidence intervals for all pairwise differences between level means using Tukey's method (also called Tukey-Kramer in the unbalanced case), and then enter a family error rate between 0.5 and 0.001. Values greater than or equal to 1.0 are interpreted as percentages. The default error rate is 0.05.

Fisher's, individual error rate: Check to obtain confidence intervals for all pairwise differences between level means using Fisher's LSD procedure, and then enter an individual rate between 0.5 and 0.001. Values greater than or equal to 1.0 are interpreted as percentages. The default error rate is 0.05.

Dunnett's family error rate: Check to obtain a two-sided confidence interval for the difference between each treatment mean and a control mean, and then enter a family error rate between 0.5 and 0.001. Values greater than or equal to 1.0 are interpreted as percentages. The default error rate is 0.05.

Control group level: Enter the value for the control group factor level. (IMPORTANT: For text variables, you must enclose factor levels in double quotes, even if there are no spaces in them.)

Hsu's MCB, family error rate: Check to obtain a confidence interval for the difference between each level mean and the best of the other level means [9]. There are two choices for "best." If the smallest mean is considered the best, set K = −1; if the largest is considered the best, set K = 1. Specify a family error rate between 0.5 and 0.001. Values greater than or equal to 1.0 are interpreted as percentages. The default error rate is 0.05.

Largest is best: Choose to have the largest mean considered the best.

Smallest is best: Choose to have the smallest mean considered the best.

One-Way Analysis of Variance − Graphs

Stat > ANOVA > One-way > Graphs

Displays an individual value plot, a boxplot, and residual plots. You do not have to store the residuals in order to produce the residual plots.

Dialog box items

Individual value plot: Check to display an individual value plot of each sample.

Boxplots of data: Check to display a boxplot of each sample.

Residual Plots

Individual plots: Choose to display one or more plots.

Histogram of residuals: Check to display a histogram of the residuals.

Normal plot of residuals: Check to display a normal probability plot of the residuals.

Residuals versus fits: Check to plot the residuals versus the fitted values.

Residuals versus order: Check to plot the residuals versus the order of the data. The row number for each data point is shown on the x-axis−for example, 1 2 3 4... n.

Four in one: Choose to display a layout of a histogram of residuals, normal plot of residuals, plot of residuals versus fits, and plot of residuals versus order.

Residuals versus the variables: Enter one or more columns containing the variables against which you want to plot the residuals. Minitab displays a separate graph for each column.

Example of a one-way analysis of variance with multiple comparisons

You design an experiment to assess the durability of four experimental carpet products. You place a sample of each of the carpet products in four homes and you measure durability after 60 days. Because you wish to test the equality of means and to assess the differences in means, you use the one-way ANOVA procedure (data in stacked form) with multiple comparisons. Generally, you would choose one multiple comparison method as appropriate for your data. However, two methods are selected here to demonstrate Minitab's capabilities.

1 Open the worksheet EXH_AOV.MTW.
2 Choose **Stat > ANOVA > One-Way**.
3 In **Response**, enter *Durability*. In **Factor**, enter *Carpet*.
4 Click **Comparisons**. Check **Tukey's, family error rate**. Check **Hsu's MCB, family error rate** and enter *10*.
5 Click **OK** in each dialog box.

Session window output

One-way ANOVA: Durability versus Carpet

```
Source  DF     SS    MS     F      P
Carpet   3  146.4  48.8  3.58  0.047
Error   12  163.5  13.6
Total   15  309.9

S = 3.691   R-Sq = 47.24%   R-Sq(adj) = 34.05%


                         Individual 95% CIs For Mean Based on
                         Pooled StDev
Level  N    Mean  StDev  ---------+---------+---------+---------+
1      4  14.483  3.157            (-------*-------)
2      4   9.735  3.566  (-------*--------)
3      4  12.808  1.506         (-------*-------)
4      4  18.115  5.435                   (-------*-------)
                         ---------+---------+---------+---------+
                              10.0      15.0      20.0      25.0

Pooled StDev = 3.691


Hsu's MCB (Multiple Comparisons with the Best)

Family error rate = 0.1
Critical value = 1.87

Intervals for level mean minus largest of other level means

Level    Lower  Center  Upper    --+---------+---------+---------+-------
1       -8.511  -3.632  1.246           (-------*-------)
2      -13.258  -8.380  0.000    (-------*-------------)
3      -10.186  -5.308  0.000         (-------*--------)
4       -1.246   3.632  8.511                        (-------*-------)
                                 --+---------+---------+---------+-------
                               -12.0      -6.0       0.0       6.0


Tukey 95% Simultaneous Confidence Intervals
All Pairwise Comparisons among Levels of Carpet

Individual confidence level = 98.83%


Carpet = 1 subtracted from:

Carpet    Lower  Center   Upper  ------+---------+---------+---------+---
2       -12.498  -4.748   3.003      (------*-------)
3        -9.426  -1.675   6.076         (------*-------)
4        -4.118   3.632  11.383              (-------*------)

                                 ------+---------+---------+---------+---
                                     -10         0        10        20


Carpet = 2 subtracted from:

Carpet   Lower  Center   Upper  ------+---------+---------+---------+---
3       -4.678   3.073  10.823            (-------*-------)
4        0.629   8.380  16.131                  (------*-------)
                                ------+---------+---------+---------+---
                                    -10         0        10        20


Carpet = 3 subtracted from:

Carpet   Lower  Center   Upper  ------+---------+---------+---------+---
```

```
4       -2.443   5.308  13.058                (------*-------)
                                    ------+---------+---------+---------+---
                                        -10         0        10        20
```

Interpreting the results

In the ANOVA table, the p-value (0.047) for Carpet indicates that there is sufficient evidence that not all the means are equal when alpha is set at 0.05. To explore the differences among the means, examine the multiple comparison results.

Hsu's MCB comparisons

Hsu's MCB (Multiple Comparisons with the Best) compares each mean with the best (largest) of the other means. Minitab compares the means of carpets 1, 2, and 3 to the carpet 4 mean because it is the largest. Carpet 1 or 4 may be best because the corresponding confidence intervals contain positive values. No evidence exists that carpet 2 or 3 is the best because the upper interval endpoints are 0, the smallest possible value.

Note You can describe the potential advantage or disadvantage of any of the contenders for the best by examining the upper and lower confidence intervals. For example, if carpet 1 is best, it is no more than 1.246 better than its closest competitor, and it may be as much as 8.511 worse than the best of the other level means.

Tukey's comparisons

Tukey's test provides 3 sets of multiple comparison confidence intervals:

- Carpet 1 mean subtracted from the carpet 2, 3, and 4 means: The first interval in the first set of the Tukey's output (−12.498, −4.748, 3.003) gives the confidence interval for the carpet 1 mean subtracted from the carpet 2 mean. You can easily find confidence intervals for entries not included in the output by reversing both the order and the sign of the interval values. For example, the confidence interval for the mean of carpet 1 minus the mean of carpet 2 is (−3.003, 4.748, 12.498). For this set of comparisons, none of the means are statistically different because all of the confidence intervals include 0.
- Carpet 2 mean subtracted from the carpet 3 and 4 means: The means for carpets 2 and 4 are statistically different because the confidence interval for this combination of means (0.629, 8.380, 16.131) excludes zero.
- Carpet 3 mean subtracted from the carpet 4 mean: Carpets 3 and 4 are not statistically different because the confidence interval includes 0.

By not conditioning upon the F-test, differences in treatment means appear to have occurred at family error rates of 0.10. If Hsu's MCB method is a good choice for these data, carpets 2 and 3 might be eliminated as a choice for the best. When you use Tukey's method, the mean durability for carpets 2 and 4 appears to be different.

One-Way (Unstacked)

One-way and two-way ANOVA models

- One-way analysis of variance tests the equality of population means when classification is by one variable. The classification variable, or factor, usually has three or more levels (one-way ANOVA with two levels is equivalent to a t-test), where the **level** represents the treatment applied. For example, if you conduct an experiment where you measure durability of a product made by one of three methods, these methods constitute the levels. The one-way procedure also allows you to examine differences among means using multiple comparisons.
- Two-way analysis of variance performs an analysis of variance for testing the equality of populations means when classification of treatments is by two variables or factors. In two-way ANOVA, the data must be balanced (all cells must have the same number of observations) and factors must be fixed.

  If you wish to specify certain factors to be random, use Balanced ANOVA if your data are balanced; use General Linear Models if your data are unbalanced or if you wish to compare means using multiple comparisons.

One-Way Analysis of Variance (Unstacked)

Stat > ANOVA > One-Way (Unstacked)

Performs a one-way analysis of variance, with each group in a separate column. If your response data are stacked in one column with another column of level values identifying the population, use Stat > ANOVA > One-Way.

You can also perform multiple comparisons and display graphs of your data.

Dialog box items

Responses [in separate columns]: Enter the columns containing the separate response variables.

Store residuals: Check to store residuals in the next available columns. The number of residual columns will match the number of response columns.

Store fits: Check to store the fitted values in the next available column.

Confidence level: Enter the confidence level. For example, enter 90 for 90%. The default is 95%.

<Comparisons>

<Graphs>

Data – One-Way (Unstacked)

The response variable must be numeric. Enter the sample data from each population into separate columns of your worksheet.

Note If your response data are stacked in one column with another column of level values identifying the population, use Stat > ANOVA > One-Way.

To perform a one-way analysis of variance with unstacked data

1 Choose **Stat > ANOVA > One-Way (Unstacked)**.

2 In **Responses (in separate columns)**, enter the columns containing the separate response variables.

3 If you like, use any dialog box options, then click **OK**.

Comparisons – One-Way Multiple Comparisons with Unstacked Data

Stat > ANOVA > One-Way (Unstacked) > Comparisons

Use to generate confidence intervals for the differences between means, using four different methods: Tukey's, Fisher's, Dunnett's, and Hsu's MCB. Tukey's and Fisher's methods provide confidence intervals for all pairwise differences between level means. Dunnett's method provides a confidence interval for the difference between each treatment mean and a control mean. Hsu's MCB method provides a confidence interval for the difference between each level mean and the best of the other level means. Tukey's, Dunnett's, and Hsu's MCB tests use a family error rate, whereas Fisher's LSD procedure uses an individual error rate.

Which multiple comparison test to use depends on the desired inference. Using Tukey's all-pairwise approach is inefficient when Dunnett's or Hsu's MCB is suitable, because Tukey's confidence intervals are wider and the hypothesis tests less powerful for a given family error rate. For the same reasons, Hsu's MCB is superior to Dunnett's if you want to eliminate levels that are not the best and to identify those that are best or close to the best. The choice of Tukey's versus Fisher's depends on which error rate, family or individual, you wish to specify.

Dialog box items

Tukey's, family error rate: Check to obtain confidence intervals for all pairwise differences between level means using Tukey's method (also called Tukey-Kramer in the unbalanced case), then enter a family error rate between 0.5 and 0.001. Values greater than or equal to 1.0 are interpreted as percentages. The default error rate is 0.05.

Fisher's, individual error rate: Check to obtain confidence intervals for all pairwise differences between level means using Fisher's LSD procedure, then enter an individual rate between 0.5 and 0.001. Values greater than or equal to 1.0 are interpreted as percentages. The default error rate is 0.05.

Dunnett's family error rate: Check to obtain a two-sided confidence interval for the difference between each treatment mean and a control mean, then enter a family error rate between 0.5 and 0.001. Values greater than or equal to 1.0 are interpreted as percentages. The default error rate is 0.05.

Control group level: Enter the column with the control group data.

Hsu's MCB, family error rate: Check to obtain a confidence interval for the difference between each level mean and the best of the other level means [9]. There are two choices for "best." If the smallest mean is considered the best, set K = –1; if the largest is considered the best, set K = 1. Specify a family error rate between 0.5 and 0.001. Values greater than or equal to 1.0 are interpreted as percentages. The default error rate is 0.05.

Largest is best: Choose to have the largest mean considered the best.

Smallest is best: Choose to have the smallest mean considered the best.

One-Way Analysis of Variance – Graphs

Stat > ANOVA > One-Way (Unstacked) > Graphs

Displays individual value plots and boxplots for each sample and residual plots.

Dialog box items

Individual value plot: Check to display an individual value plot of each sample. The sample mean is shown on each dotplot.

Boxplots of data: Check to display a boxplot of each sample. The sample mean is shown on each boxplot.

Residual Plots

Individual plots: Choose to display one or more plots.

Histogram of residuals: Check to display a histogram of the residuals.

Normal plot of residuals: Check to display a normal probability plot of the residuals.

Residuals versus fits: Check to plot the residuals versus the fitted values.

Three in one: Choose to display a layout of a histogram of residuals, normal plot of residuals, and a plot of residuals versus fits.

Discussion of Multiple Comparisons

The multiple comparisons are presented as a set of confidence intervals, rather than as a set of hypothesis tests. This allows you to assess the practical significance of differences among means, in addition to statistical significance. As usual, the null hypothesis of no difference between means is rejected if and only if zero is not contained in the confidence interval.

The selection of the appropriate multiple comparison method depends on the desired inference. It is inefficient to use the Tukey all-pairwise approach when Dunnett or MCB is suitable, because the Tukey confidence intervals will be wider and the hypothesis tests less powerful for a given family error rate. For the same reasons, MCB is superior to Dunnett if you want to eliminate levels that are not the best and to identify those that are best or close to the best. The choice of Tukey versus Fisher methods depends on which error rate, family or individual, you wish to specify.

Individual error rates are exact in all cases. Family error rates are exact for equal group sizes. If group sizes are unequal, the true family error rate for Tukey, Fisher, and MCB will be slightly smaller than stated, resulting in conservative confidence intervals [4,22]. The Dunnett family error rates are exact for unequal sample sizes.

The results of the one-way F-test and multiple comparisons can conflict. For example, it is possible for the F-test to reject the null hypothesis of no differences among the level means, and yet all the Tukey pairwise confidence intervals contain zero. Conversely, it is possible for the F-test to fail to reject, and yet have one or more of the Tukey pairwise confidence intervals not include zero. The F-test has been used to protect against the occurrence of false positive differences in means. However, Tukey, Dunnett, and MCB have protection against false positives built in, while Fisher only benefits from this protection when all means are equal. If the use of multiple comparisons is conditioned upon the significance of the F-test, the error rate can be higher than the error rate in the unconditioned application of multiple comparisons [15].

Two-Way

One-way and two-way ANOVA models

- One-way analysis of variance tests the equality of population means when classification is by one variable. The classification variable, or factor, usually has three or more levels (one-way ANOVA with two levels is equivalent to a t-test), where the **level** represents the treatment applied. For example, if you conduct an experiment where you measure durability of a product made by one of three methods, these methods constitute the levels. The one-way procedure also allows you to examine differences among means using multiple comparisons.
- Two-way analysis of variance performs an analysis of variance for testing the equality of populations means when classification of treatments is by two variables or factors. In two-way ANOVA, the data must be balanced (all cells must have the same number of observations) and factors must be fixed.

  If you wish to specify certain factors to be random, use Balanced ANOVA if your data are balanced; use General Linear Models if your data are unbalanced or if you wish to compare means using multiple comparisons.

Two-Way Analysis of Variance

Stat > ANOVA > Two-Way

A two-way analysis of variance tests the equality of populations means when classification of treatments is by two variables or factors. For this procedure, the data must be balanced (all cells must have the same number of observations) and factors must be fixed.

To display cell means and standard deviations, use Cross Tabulation and Chi-Square.

If you wish to specify certain factors to be random, use Balanced ANOVA if your data are balanced. Use General Linear Model if your data are unbalanced or if you wish to compare means using multiple comparisons.

Dialog box items

Response: Enter the column containing the response variable.

Row Factor: Enter one of the factor level columns.

Display means: Check to compute marginal means and confidence intervals for each level of the row factor.

Column factor: Enter the other factor level column.

Display means: Check to compute marginal means and confidence intervals for each level of the column factor.

Store residuals: Check to store the residuals.

Store fits: Check to store the fitted value for each group.

Confidence level: Enter the level for the confidence intervals for the individual means. For example, enter 90 for 90%. The default is 95%.

Fit additive model: Check to fit a model without an interaction term. In this case, the fitted value for cell (i,j) is (mean of observations in row i) + (mean of observations in row j) – (mean of all observations).

<Graphs>

Data – Two-Way

The response variable must be numeric and in one worksheet column. You must have a single factor level column for each of the two factors. These can be numeric, text, or date/time. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. See Ordering Text Categories. You must have a balanced design (same number of observations in each treatment combination) with fixed factors. You can use Make Patterned Data to enter repeated factor levels.

To perform a two-way analysis of variance

1 Choose **Stat > ANOVA > Two-Way**.

2 In **Response**, enter the column containing the response variable.

3 In **Row Factor**, enter one of the factor level columns.

4 In **Column Factor**, enter the other factor level column.

5 If you like, use any dialog box options, then click **OK**.

Two-Way Analysis of Variance – Graphs

Stat > ANOVA > Two-Way > Graphs

Displays residual plots. You do not have to store the residuals in order to produce these plots.

Dialog box items

Individual value plot: Check to display an individual value plot of each sample.

Boxplots of data: Check to display a boxplot of each sample.

Residual Plots

Individual plots: Choose to display one or more plots.

Histogram of residuals: Check to display a histogram of the residuals.

Normal plot of residuals: Check to display a normal probability plot of the residuals.

Residuals versus fits: Check to plot the residuals versus the fitted values.

Residuals versus order: Check to plot the residuals versus the order of the data. The row number for each data point is shown on the x-axis – for example, 1 2 3 4... n.

Four in one: Choose to display a layout of a histogram of residuals, a normal plot of residuals, a plot of residuals versus fits, and a plot of residuals versus order.

Residuals versus the variables: Enter one or more columns containing the variables against which you want to plot the residuals. Minitab displays a separate graph for each column.

Example of a Two-Way Analysis of Variance

You as a biologist are studying how zooplankton live in two lakes. You set up twelve tanks in your laboratory, six each with water from one of the two lakes. You add one of three nutrient supplements to each tank and after 30 days you count the zooplankton in a unit volume of water. You use two-way ANOVA to test if the population means are equal, or equivalently, to test whether there is significant evidence of interactions and main effects.

1 Open the worksheet EXH_AOV.MTW.
2 Choose **Stat > ANOVA > Two-Way**.
3 In **Response**, enter *Zooplankton*.
4 In **Row factor**, enter *Supplement*. Check **Display means**.
5 In **Column factor**, enter *Lake*. Check **Display means**. Click **OK**.

Session window output

Two-way ANOVA: Zooplankton versus Supplement, Lake

```
Source       DF       SS       MS     F      P
Supplement    2  1918.50  959.250  9.25  0.015
Lake          1    21.33   21.333  0.21  0.666
Interaction   2   561.17  280.583  2.71  0.145
Error         6   622.00  103.667
Total        11  3123.00

S = 10.18   R-Sq = 80.08%   R-Sq(adj) = 63.49%


                   Individual 95% CIs For Mean Based on
                   Pooled StDev
Supplement   Mean  --+---------+---------+---------+-------
1           43.50     (-------*-------)
2           68.25                     (--------*-------)
3           39.75  (--------*-------)
                   --+---------+---------+---------+-------
                    30        45        60        75


                   Individual 95% CIs For Mean Based on
                   Pooled StDev
Lake         Mean  -----+---------+---------+---------+----
Dennison  51.8333      (----------------*----------------)
Rose      49.1667  (----------------*----------------)
                   -----+---------+---------+---------+----
                     42.0      48.0      54.0      60.0
```

Interpreting the results

The default output for two-way ANOVA is the analysis of variance table. For the zooplankton data, there is no significant evidence for a supplement∗lake interaction effect or a lake main effect if your acceptable a value is less than 0.145 (the p-value for the interaction F-test). There is significant evidence for supplement main effects, as the F-test p-value is 0.015.

As requested, the means are displayed with individual 95% confidence intervals. Supplement 2 appears to have provided superior plankton growth in this experiment. These are t-distribution confidence intervals calculated using the error degrees of freedom and the pooled standard deviation (square root of the mean square error). If you want to examine simultaneous differences among means using multiple comparisons, use General Linear Model.

Analysis of Means

Overview of Analysis of Means

Analysis of Means (ANOM) is a graphical analog to ANOVA, and tests the equality of population means. ANOM [16] was developed to test main effects from a designed experiment in which all factors are fixed. This procedure is used for one-factor designs. Minitab uses an extension of ANOM or ANalysis Of Mean treatment Effects (ANOME) [24] to test the significance of mean treatment effects for two-factor designs.

An ANOM chart can be described in two ways: by its appearance and by its function. In appearance, it resembles a Shewhart control chart. In function, it is similar to ANOVA for detecting differences in population means [13]. The null hypotheses for ANOM and ANOVA are the same: both methods test for a lack of homogeneity among means. However, the alternative hypotheses are different [16]. The alternative hypothesis for ANOM is that one of the population means is different from the other means, which are equal. The alternative hypothesis for ANOVA is that the variability among population means is greater than zero.

For most cases, ANOVA and ANOM will likely give similar results. However, there are some scenarios where the two methods might be expected to differ:

- If one group of means is above the grand mean and another group of means is below the grand mean, the F-test for ANOVA might indicate evidence for differences where ANOM might not.
- If the mean of one group is separated from the other means, the ANOVA F-test might not indicate evidence for differences whereas ANOM might flag this group as being different from the grand mean.

Refer to [21], [22], [23], and [24] for an introduction to the analysis of means.

ANOM can be used if you assume that the response follows a normal distribution, similar to ANOVA, and the design is one-way or two-way. You can also use ANOM when the response follows either a binomial distribution or a Poisson distribution.

Analysis of Means

Stat > ANOVA > Analysis of Means

Draws an Analysis of Means chart (ANOM) for normal, binomial, and Poisson data and optionally prints a summary table for normal and binomial data.

Dialog box items

Response: Enter the column containing the response variable. The meaning of and limitations for the response variable vary depending on whether your data follow a normal, binomial, or Poisson distribution. See Data for Analysis of Means for more information.

Distribution of Data:

Normal: Choose if the response data follow a normal distribution (measurement data).

Factor 1: Enter the column containing the levels for the first factor. If you enter a single factor, Analysis of Means produces a single plot showing the means for each level of the factor.

Factor 2 (optional): Enter the column containing the levels for the second factor. If you enter two factors, Analysis of Means produces three plots−one showing the interaction effects, one showing the main effects for the first factor, and one showing the main effects for the second factor.

Binomial: Choose if the response data follow a binomial distribution. The sample size must be constant (balanced design), and the sample size must be large enough to ensure that the normal approximation to the binomial is valid. This usually implies that $np > 5$ and $n(1-p) > 5$, where p is the proportion of defects.

Sample size: Enter a number or a stored constant to specify the sample size.

Poisson: Choose if the response data follow a Poisson distribution. The Poisson distribution can be adequately approximated by a normal distribution if the mean of the Poisson distribution is at least five. Therefore, when the Poisson mean is large enough, you can test for equality of population means using this procedure.

Alpha level: Enter a value for the error rate, or alpha-level. The number you enter must be between 0 and 1. The decision lines on the ANOM chart are based on an experiment-wide error rate, similar to what you might use when making pairwise comparisons or contrasts in an ANOVA.

Title: Type a new title to replace the plot's default title.

Data − Analysis of Means (normal)

Your response data may be numeric or date/time and must be entered into one column. Factor columns may be numeric, text, or date/time and may contain any values. The response and factor columns must be of the same length. Minitab's capability to enter patterned data can be helpful in entering numeric factor levels; see Make Patterned Data to enter repeated factor levels. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order; see Ordering Text Categories.

One-way designs may be balanced or unbalanced and can have up to 100 levels. Two-way designs must be balanced and can have up to 50 levels for each factor. All factors must be fixed.

Rows with missing data are automatically omitted from calculations. If you have two factors, the design must be balanced after omitting rows with missing values.

Data − Analysis of Means (binomial)

The response data are the numbers of defectives (or defects) found in each sample, with a maximum of 500 samples. You must enter these data in one column.

Because the decision limits in the ANOM chart are based upon the normal distribution, one of the assumptions that must be met when the response data are binomial is that the sample size is large enough to ensure that the normal approximation to the binomial is valid. A general rule of thumb is to only use ANOM if $np > 5$ and $n(1 - p) > 5$, where n is the sample size and p is the proportion of defectives. The second assumption is that all of the samples are the same size. See [24] for more details.

A sample with a missing response value (∗) is automatically omitted from the analysis.

Data – Analysis of Means (Poisson)

The response data are the numbers of defects found in each sample. You can have up to 500 samples.

The Poisson distribution can be adequately approximated by a normal distribution if the mean of the Poisson distribution is at least 5. When the Poisson mean is large enough, you can apply analysis of means to data from a Poisson distribution to test if the population means are equal to the grand mean.

A sample with a missing response value (*) is automatically omitted from the analysis.

To perform an analysis of means

1 Choose **Stat > ANOVA > Analysis of Means**.

2 In **Response**, enter a numeric column containing the response variable.

3 Under **Distribution of Data**, choose **Normal**, **Binomial**, or **Poisson**.

- If you choose **Normal**, you can analyze either a one-way or two-way design. For a one-way design, enter the column containing the factor levels in **Factor 1**. For a two-way design, enter the columns containing the factor levels in **Factor 1** and **Factor 2**.
- If you choose **Binomial**, type a number in **Sample size**.

4 If you like, use one or more of the dialog box options, then click **OK**.

Example of a two-way analysis of means

You perform an experiment to assess the effect of three process time levels and three strength levels on density. You use analysis of means for normal data and a two-way design to identify any significant interactions or main effects.

1 Open the worksheet EXH_AOV.MTW.

2 Choose **Stat > ANOVA > Analysis of Means**.

3 In **Response**, enter *Density*.

4 Choose **Normal**.

5 In **Factor 1**, enter *Minutes*. In **Factor 2**, enter *Strength*. Click **OK**.

Graph window output

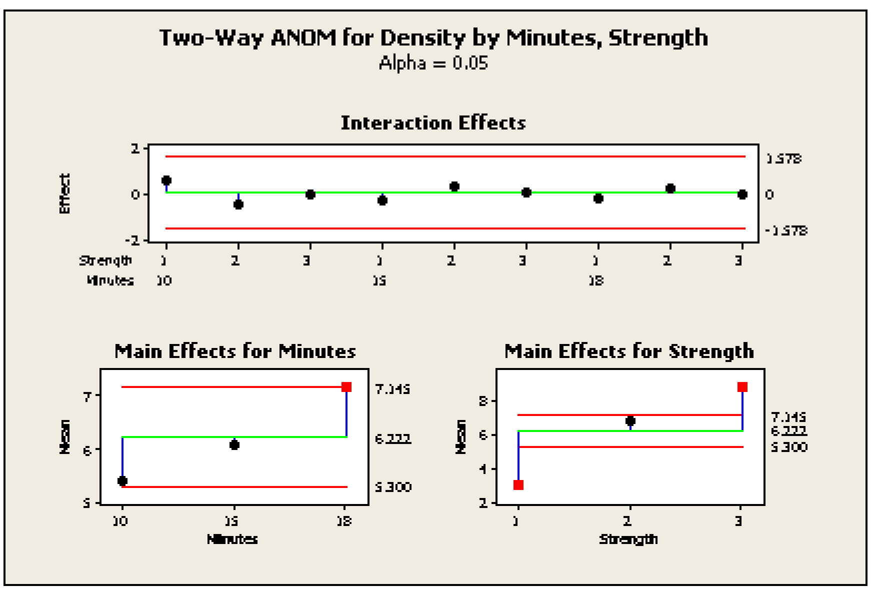

Interpreting the results

Minitab displays three plots with a two-way ANOM to show the interaction effects, main effects for the first factor, and main effects for the second factor. ANOM plots have a center line and decision limits. If a point falls outside the decision limits, then there is significant evidence that the mean represented by that point is different from the grand mean. With a two-way ANOM, look at the interaction effects first. If there is significant evidence for interaction, it usually does not make sense to consider main effects, because the effect of one factor depends upon the level of the other.

In this example, the interaction effects are well within the decision limits, signifying no evidence of interaction. Now you can look at the main effects. The lower two plots show the means for the levels of the two factors, with the main effect being the difference between the mean and the center line. The point representing the level 3 mean of the factor Minutes is displayed by a red asterisk, which indicates that there is significant evidence that the level 3 mean is different from the grand mean at $\alpha = 0.05$. You may wish to investigate any point near or above the decision limits. The main effects for levels 1 and 3 of factor Strength are well outside the decision limits of the lower left plot, signifying that there is evidence that these means are different from the grand mean at $\alpha = 0.05$.

Example of analysis of means for binomial response data

You count the number of rejected welds from samples of size 80 in order to identify samples whose proportions of rejects are out of line with the other samples. Because the data are binomial (two possible outcomes, constant proportion of success, and independent samples) you use analysis of means for binomial data.

1 Open the worksheet EXH_AOV.MTW.

2 Choose **Stat > ANOVA > Analysis of Means**.

3 In **Response**, enter *WeldRejects*.

4 Choose **Binomial** and type *80* in **Sample size**. Click **OK.**

Graph window output

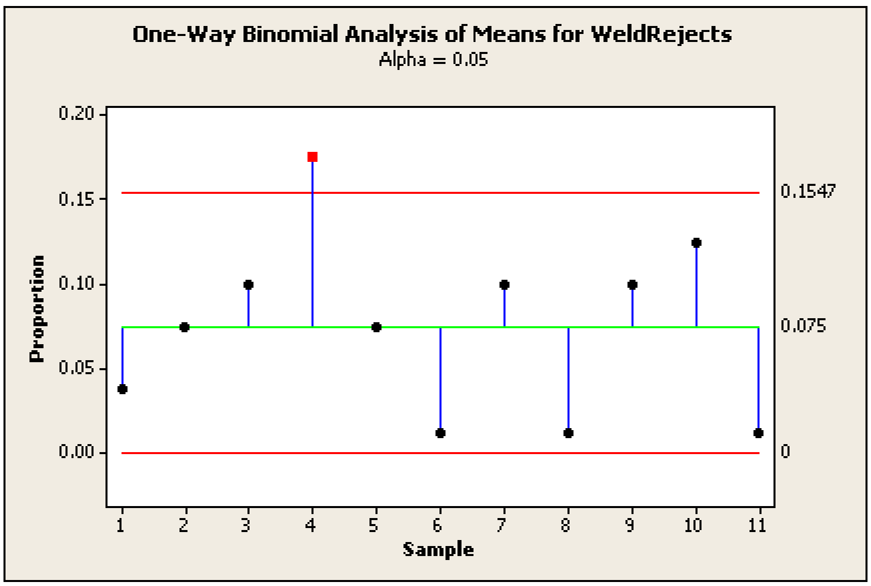

Interpreting the results

The plot displays the proportion of defects for each sample, a center line representing the average proportion, and upper and lower decision limits. If the point representing a sample falls outside the decision limits, there is significant evidence that the sample mean is different from the average

In this example, the proportion of defective welds in sample four is identified as unusually high because the point representing this sample falls outside the decision limits.

Example of analysis of means for Poisson response data

As production manager of a toy manufacturing company, you want to monitor the number of defects per sample of motorized toy cars. You monitor 20 samples of toy cars and create an analysis of means chart to examine the number of defects in each sample.

1 Open the worksheet TOYS.MTW.

2 Choose **Stat > ANOVA > Analysis of Means**.

3 In **Response**, enter *Defects*.

4 Choose **Poisson**, then click **OK**.

Graph window output

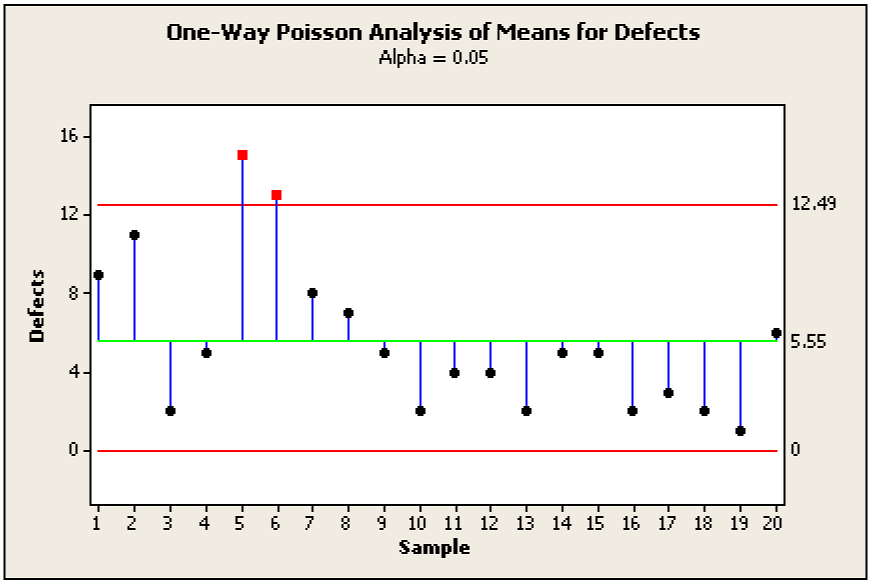

Interpreting the results

The plot displays the number of defects for each sample, a center line representing the average number of defects, and upper and lower decision limits. If the point representing a sample falls outside the decision limits, there is significant evidence exists that the sample mean is different from the average

In this example, the number of defective motorized toy cars in samples five and six is identified as being unusually high because the points representing these samples fall outside the decision limits.

Balanced ANOVA

Overview of Balanced ANOVA and GLM

Balanced ANOVA and general linear model (GLM) are ANOVA procedures for analyzing data collected with many different experimental designs. Your choice between these procedures depends upon the experimental design and the available options. The experimental design refers to the selection of units or subjects to measure, the assignment of treatments to these units or subjects, and the sequence of measurements taken on the units or subjects. Both procedures can fit univariate models to balanced data with up to 31 factors. Here are some of the other options:

	Balanced ANOVA	**GLM**
Can fit unbalanced data	no	yes
Can specify factors as random and obtain expected means squares	yes	yes
Fits covariates	no	yes
Performs multiple comparisons	no	yes
Fits restricted/unrestricted forms of mixed model	yes	unrestricted only

You can use balanced ANOVA to analyze data from balanced designs. See Balanced designs. You can use GLM to analyze data from any balanced design, though you cannot choose to fit the restricted case of the mixed model, which only balanced ANOVA can fit. See Restricted and unrestricted form of mixed models.

To classify your variables, determine if your factors are:

- crossed or nested
- fixed or random
- covariates

For information on how to specify the model, see Specifying the model terms, Specifying terms involving covariates, Specifying reduced models, and Specifying models for some specialized designs.

For easy entering of repeated factor levels into your worksheet, see Using patterned data to set up factor levels.

Balanced Analysis of Variance

Stat > ANOVA > Balanced ANOVA

Use Balanced ANOVA to perform univariate analysis of variance for each response variable.

Your design must be balanced, with the exception of one-way designs. *Balanced* means that all treatment combinations (cells) must have the same number of observations. See Balanced designs. Use General Linear Model to analyze balanced and unbalanced designs.

Factors may be crossed or nested, fixed or random. You may include up to 50 response variables and up to 31 factors at one time.

Dialog box items

Responses: Enter the columns containing the response variables.

Model: Enter the terms to be included in the model. See Specifying a Model for more information.

Random factors: Enter any columns containing random factors. Do not include model terms that involve other factors.

<Options>

<Graphs>

<Results>

<Storage>

Data – Balanced ANOVA

You need one column for each response variable and one column for each factor, with each row representing an observation. Regardless of whether factors are crossed or nested, use the same form for the data. Factor columns may be numeric, text, or date/time. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. See Ordering Text Categories. You may include up to 50 response variables and up to 31 factors at one time.

Balanced data are required except for one-way designs. The requirement for balanced data extends to nested factors as well. Suppose A has 3 levels, and B is nested within A. If B has 4 levels within the first level of A, B must have 4 levels within the second and third levels of A. Minitab will tell you if you have unbalanced nesting. In addition, the subscripts used to indicate the 4 levels of B within each level of A must be the same. Thus, the four levels of B cannot be (1 2 3 4) in level 1 of A, (5 6 7 8) in level 2 of A, and (9 10 11 12) in level 3 of A.

If any response or factor column specified contains missing data, that entire observation (row) is excluded from all computations. The requirement that data be balanced must be preserved after missing data are omitted. If an observation is missing for one response variable, that row is eliminated for all responses. If you want to eliminate missing rows separately for each response, perform a separate ANOVA for each response.

To perform a balanced ANOVA

1. Choose **Stat > ANOVA > Balanced ANOVA**.
2. In **Responses**, enter up to 50 numeric columns containing the response variables.
3. In **Model**, type the model terms you want to fit. See Specifying the Model Terms.
4. If you like, use any dialog box options, then click **OK**.

Balanced designs

Your design must be balanced to use balanced ANOVA, with the exception of a one-way design. A *balanced* design is one with equal numbers of observations at each combination of your treatment levels. A quick test to see whether or not you have a balanced design is to use **Stat > Tables > Cross Tabulation and Chi-Square**. Enter your classification variables and see if you have equal numbers of observations in each cell, indicating balanced data.

Restricted and unrestricted form of mixed models

A mixed model is one with both fixed and random factors. There are two forms of this model: one requires the crossed, mixed terms to sum to zero over subscripts corresponding to fixed effects (this is called the restricted model), and the other does not. See Example of both restricted and unrestricted forms of the mixed model. Many textbooks use the restricted model. Most statistics programs use the unrestricted model. Minitab fits the unrestricted model by default, but you can choose to fit the restricted form. The reasons to choose one form over the other have not been clearly defined in

the statistical literature. Searle et al. [25] say "that question really has no definitive, universally acceptable answer," but also say that one "can decide which is more appropriate to the data at hand," without giving guidance on how to do so.

Your choice of model form does not affect the sums of squares, degrees of freedom, mean squares, or marginal and cell means. It does affect the expected mean squares, error terms for F-tests, and the estimated variance components. See Example of both restricted and unrestricted forms of the mixed model.

Specifying a model

Specify a model in the Model text box using the form **Y = expression**. The **Y** is not included in the Model text box. The Calc > Make Patterned Data > Simple Set of Numbers command can be helpful in entering the level numbers of a factor.

Rules for Expression Models

1 ∗ indicates an interaction term. For example, A∗B is the interaction of the factors A and B.

2 ( ) indicate nesting. When B is nested within A, type B(A). When C is nested within both A and B, type C(A B). Terms in parentheses are always factors in the model and are listed with blanks between them.

3 Abbreviate a model using a | or ! to indicate crossed factors and a – to remove terms.

Models with many terms take a long time to compute.

Examples of what to type in the Model text box

Two factors crossed: `A B A*B`

Three factors crossed: `A B C A*B A*C B*C A*B*C`

Three factors nested: `A B(A) C(A B)`

Crossed and nested (B nested within A, and both crossed with C): `A B(A) C A*C B*C(A)`

When a term contains both crossing and nesting, put the ∗ (or crossed factor) first, as in C∗B(A), not B(A)∗C

Example of entering level numbers for a data set

Here is an easy way to enter the level numbers for a three-way crossed design with a, b, and c levels of factors A, B, C, with n observations per cell:

1 Choose **Calc > Make Patterned Data > Simple Set of Numbers**, and press <F3> to reset defaults. Enter *A* in **Store patterned data in**. Enter *1* in **From first value**. Enter the number of levels in A in **To last value**. Enter the product of bcn in **List the whole sequence**. Click **OK**.

2 Choose **Calc > Make Patterned Data > Simple Set of Numbers**, and press <F3> to reset defaults. Enter *B* in **Store patterned data in**. Enter *1* in **From first value**. Enter the number of levels in B in **To last value**. Enter the number of levels in A in **List each value**. Enter the product of cn in **List the whole sequence**. Click **OK**.

3 Choose **Calc > Make Patterned Data > Simple Set of Numbers**, and press <F3> to reset defaults. Enter *C* in **Store patterned data in**. Enter *1* in *From first value*. Enter the number of levels in C in *To last value*. Enter the product of ab in **List each value**. Enter the sample size n in **List the whole sequence**. Click **OK**.

Specifying the model terms

You must specify the model terms in the **Model** box. This is an abbreviated form of the statistical model that you may see in textbooks. Because you enter the response variables in **Responses,** in **Model** you enter only the variables or products of variables that correspond to terms in the statistical model. Minitab uses a simplified version of a statistical model as it appears in many textbooks. Here are some examples of statistical models and the terms to enter in **Model**. A, B, and C represent factors.

Case	Statistical model	Terms in model
Factors A, B crossed	$y_{ijk} = \mu + a_i + b_j + ab_{ij} + e_{k(ij)}$	A B A∗ B
Factors A, B, C crossed	$y_{ijkl} = \mu + a_i + b_j + c_k + ab_{ij} + ac_{ik} + bc_{jk} + abc_{ijk} + e_{l(ijk)}$	A B C A∗B A∗ C B∗C A∗ B∗C
3 factors nested (B within A, C within A and B)	$y_{ijkl} = \mu + a_i + b_{j(i)} + c_{k(ij)} + e_{l(ijk)}$	A B(A) C(AB)
Crossed and nested (B nested within A, both crossed with C)	$y_{ijkl} = \mu + a_i + b_{j(i)} + c_k + ac_{ik} + bc_{jk(i)} + e_{l(ijk)}$	A B (A) C A∗C B∗C

In Minitab's models you omit the subscripts, μ, e, and +'s that appear in textbook models. An ∗ is used for an interaction term and parentheses are used for nesting. For example, when B is nested within A, you enter B (A), and when C is nested within both A and B, you enter C (A B). Enter B(A) C(B) for the case of 3 sequentially nested factors. Terms in

parentheses are always factors in the model and are listed with blanks between them. Thus, D ∗ F (A B E) is correct but D ∗ F (A ∗ B E) and D (A ∗ B ∗ C) are not. Also, one set of parentheses cannot be used inside another set. Thus, C (A B) is correct but C (A B (A)) is not. An interaction term between a nested factor and the factor it is nested within is invalid.

See Specifying terms involving covariates for details on specifying models with covariates.

Several special rules apply to naming columns. You may omit the quotes around variable names. Because of this, variable names must start with a letter and contain only letters and numbers. Alternatively, you can use C notation (C1, C2, etc.) to denote data columns. You can use special symbols in a variable name, but then you must enclose the name in single quotes.

You can specify multiple responses. In this case, a separate analysis of variance will be performed for each response.

Specifying models for some specialized designs

Some experimental designs can effectively provide information when measurements are difficult or expensive to make or can minimize the effect of unwanted variability on treatment inference. The following is a brief discussion of three commonly used designs that will show you how to specify the model terms in Minitab. To illustrate these designs, two treatment factors (A and B) and their interaction (A∗B) are considered. These designs are not restricted to two factors, however. If your design is balanced, you can use balanced ANOVA to analyze your data. Otherwise, use GLM.

Randomized block design

A *randomized block* design is a commonly used design for minimizing the effect of variability when it is associated with discrete units (e.g. location, operator, plant, batch, time). The usual case is to randomize one replication of each treatment combination within each block. There is usually no intrinsic interest in the blocks and these are considered to be random factors. The usual assumption is that the block by treatment interaction is zero and this interaction becomes the error term for testing treatment effects. If you name the block variable as Block, enter *Block A B A∗B* in **Model** and enter *Block* in **Random Factors**.

Split-plot design

A *split-plot* design is another blocking design, which you can use if you have two or more factors. You might use this design when it is more difficult to randomize one of the factors compared to the other(s). For example, in an agricultural experiment with the factors variety and harvest date, it may be easier to plant each variety in contiguous rows and to randomly assign the harvest dates to smaller sections of the rows. The block, which can be replicated, is termed the *main plot* and within these the smaller plots (variety strips in example) are called *subplots*.

This design is frequently used in industry when it is difficult to randomize the settings on machines. For example, suppose that factors are temperature and material amount, but it is difficult to change the temperature setting. If the blocking factor is operator, observations will be made at different temperatures with each operator, but the temperature setting is held constant until the experiment is run for all material amounts. In this example, the plots under operator constitute the main plots and temperatures constitute the subplots.

There is no single error term for testing all factor effects in a split-plot design. If the levels of factor A form the subplots, then the mean square for Block ∗ A will be the error term for testing factor A. There are two schools of thought for what should be the error term to use for testing B and A ∗ B. If you enter the term Block ∗ B, the expected mean squares show that the mean square for Block ∗ B is the proper term for testing factor B and that the remaining error (which is Block ∗ A ∗ B) will be used for testing A ∗ B. However, it is often assumed that the Block ∗ B and Block ∗ A ∗ B interactions do not exist and these are then lumped together into error [6]. You might also pool the two terms if the mean square for Block ∗ B is small relative to Block ∗ A ∗ B. If you don't pool, enter *Block A Block ∗ A B Block ∗B A ∗ B* in **Model** and what is labeled as Error is really Block ∗ A ∗ B. If you do pool terms, enter *Block A Block ∗ A B A ∗ B* in **Model** and what is labeled as Error is the set of pooled terms. In both cases enter *Block* in **Random Factors**.

Latin square with repeated measures design

A *repeated measures* design is a design where repeated measurements are made on the same subject. There are a number of ways in which treatments can be assigned to subjects. With living subjects especially, systematic differences (due to learning, acclimation, resistance, etc.) between successive observations may be suspected. One common way to assign treatments to subjects is to use a Latin square design. An advantage of this design for a repeated measures experiment is that it ensures a balanced fraction of a complete factorial (i.e. all treatment combinations represented) when subjects are limited and the sequence effect of treatment can be considered to be negligible.

A *Latin square* design is a blocking design with two orthogonal blocking variables. In an agricultural experiment there might be perpendicular gradients that might lead you to choose this design. For a repeated measures experiment, one blocking variable is the group of subjects and the other is time. If the treatment factor B has three levels, b_1, b_2, and b_3, then one of twelve possible Latin square randomizations of the levels of B to subjects groups over time is:

	Time 1	Time 2	Time 3
Group 1	b2	b3	b1
Group 2	b3	b1	b2
Group 3	b1	b2	b3

The subjects receive the treatment levels in the order specified across the row. In this example, group 1 subjects would receive the treatments levels in order b2, b3, b1. The interval between administering treatments should be chosen to minimize carryover effect of the previous treatment.

This design is commonly modified to provide information on one or more additional factors. If each group was assigned a different level of factor A, then information on the A and A ∗ B effects could be made available with minimal effort if an assumption about the sequence effect given to the groups can be made. If the sequence effects are negligible compared to the effects of factor A, then the group effect could be attributed to factor A. If interactions with time are negligible, then partial information on the A ∗ B interaction may be obtained [29]. In the language of repeated measures designs, factor A is called a *between-subjects* factor and factor B a *within-subjects* factor.

Let's consider how to enter the model terms into Minitab. If the group or A factor, subject, and time variables were named A, Subject, and Time, respectively, enter *A Subject(A) Time B A ∗ B* in **Model** and enter *Subject* in **Random Factors**.

It is not necessary to randomize a repeated measures experiments according to a Latin square design. See Example of a repeated measures design for a repeated measures experiment where the fixed factors are arranged in a complete factorial design.

Specifying reduced models

You can fit *reduced* models. For example, suppose you have a three factor design, with factors, A, B, and C. The *full* model would include all one factor terms: A, B, C, all two-factor interactions: A ∗ B, A ∗ C, B ∗ C, and the three-factor interaction: A ∗ □B ∗ C. It becomes a reduced model by omitting terms. You might reduce a model if terms are not significant or if you need additional error degrees of freedom and you can assume that certain terms are zero. For this example, the model with terms A B C A ∗ B is a reduced three-factor model.

One rule about specifying reduced models is that they must be hierarchical. That is, for a term to be in the model, all lower order terms contained in it must also be in the model. For example, suppose there is a model with four factors: A, B, C, and D. If the term A ∗ B ∗ C is in the model then the terms A B C A ∗ B A ∗ C B ∗ C must also be in the model, though any terms with D do not have to be in the model. The hierarchical structure applies to nesting as well. If B (A) is in the model, then A must be also.

Because models can be quite long and tedious to type, two shortcuts have been provided. A vertical bar indicates crossed factors, and a minus sign removes terms.

Long form	**Short form**
A B C A ∗ B A ∗ C B ∗ C A ∗ B ∗ C	A | B | C
A B C A ∗ B A ∗ C B ∗ C	A | B | C – A ∗ B ∗ C
A B C B ∗ C E	A B | C E
A B C D A ∗ B A ∗ C A ∗ D B ∗ C B ∗ D C ∗ D A ∗ B ∗ D A ∗ C ∗ D B ∗ C ∗ D	A | B | C | D – A ∗ B ∗ C – A ∗ B ∗ C ∗ D
A B (A) C A ∗ C B ∗ C	A | B (A) | C

In general, all crossings are done for factors separated by bars unless the cross results in an illegal term. For example, in the last example, the potential term A ∗ B (A) is illegal and Minitab automatically omits it. If a factor is nested, you must indicate this when using the vertical bar, as in the last example with the term B (A).

Using patterned data to set up factor levels

Minitab's set patterned data capability can be helpful when entering numeric factor levels. For example, to enter the level values for a three-way crossed design with a, b, and c (a, b, and c represent numbers) levels of factors A, B, C, and n observations per cell, fill out the Calc > Set Patterned Data > Simple Set of Numbers dialog box and execute 3 times, once for each factor, as shown:

		Factor	
Dialog item	**A**	**B**	**C**
From first value	1	1	1
From last value	a	b	c
List each value	bcn	cn	n
List the whole sequence	1	a	ab

Balanced ANOVA – Options

Stat > ANOVA > Balanced ANOVA > Options

Use to fit a restricted model.

Dialog box items

Use the restricted form of the model: Check to fit a restricted model, with mixed interaction terms restricted to sum to zero over the fixed effects. Minitab will fit an unrestricted model if this box is left unchecked. See Restricted and unrestricted form of mixed models.

Balanced ANOVA – Graphs

Stat > ANOVA > Balanced ANOVA > Graphs

Displays residual plots. You do not have to store the residuals in order to produce these plots.

Dialog box items

Residual Plots

Individual plots: Choose to display one or more plots.

Histogram of residuals: Check to display a histogram of the residuals.

Normal plot of residuals: Check to display a normal probability plot of the residuals.

Residuals versus fits: Check to plot the residuals versus the fitted values.

Residuals versus order: Check to plot the residuals versus the order of the data. The row number for each data point is shown on the x-axis–for example, 1 2 3 4... n.

Four in one: Choose to display a layout of a histogram of residuals, a normal plot of residuals, a plot of residuals versus fits, and a plot of residuals versus order.

Residuals versus the variables: Enter one or more columns containing the variables against which you want to plot the residuals. Minitab displays a separate graph for each column.

Balanced ANOVA – Results

Stat > ANOVA > Balanced ANOVA > Results

Use to control the Session window output.

Dialog box items

Display expected mean squares and variance components: Check to display a table that contains expected mean squares, estimated variance components, and the error term (the denominator) used in each F-test. See Expected means squares.

Display means corresponding to the terms: Enter terms for which a table of sample sizes and means will be printed. These terms must be in the model.

Expected mean squares

If you do not specify any factors to be random, Minitab will assume that they are fixed. In this case, the denominator for F-statistics will be the MSE. However, for models which include random terms, the MSE is not always the correct error term. You can examine the expected means squares to determine the error term that was used in the F-test.

When you select **Display expected mean squares and variance components** in the **Results** subdialog box, Minitab will print a table of expected mean squares, estimated variance components, and the error term (the denominator mean squares) used in each F-test. The *expected mean squares* are the expected values of these terms with the specified model. If there is no exact F-test for a term, Minitab solves for the appropriate error term in order to construct an approximate F-test. This test is called a *synthesized test*.

The estimates of variance components are the usual unbiased analysis of variance estimates. They are obtained by setting each calculated mean square equal to its expected mean square, which gives a system of linear equations in the unknown variance components that is then solved. Unfortunately, this method can result in negative estimates, which should be set to zero. Minitab, however, prints the negative estimates because they sometimes indicate that the model being fit is inappropriate for the data. Variance components are not estimated for fixed terms.

Balanced ANOVA – Storage

Stat > ANOVA > Balanced ANOVA > Storage

Stores the fitted values and residuals.

Dialog box items

Fits: Check to store the fitted values for each observation in the data set in the next available columns, using one column for each response.

Residuals: Check to store the residuals using one column for each response.

Example of ANOVA with Two Crossed Factors

An experiment was conducted to test how long it takes to use a new and an older model of calculator. Six engineers each work on both a statistical problem and an engineering problem using each calculator model and the time in minutes to solve the problem is recorded. The engineers can be considered as blocks in the experimental design. There are two factors– type of problem, and calculator model– each with two levels. Because each level of one factor occurs in combination with each level of the other factor, these factors are crossed. The example and data are from Neter, Wasserman, and Kutner [19], page 936.

1 Open the worksheet EXH_AOV.MTW.
2 Choose **Stat > ANOVA > Balanced ANOVA**.
3 In **Responses**, enter *SolveTime*.
4 In **Model**, type *Engineer ProbType* | *Calculator*.
5 In **Random Factors**, enter *Engineer*.
6 Click **Results**. In **Display means corresponding to the terms**, type *ProbType* | *Calculator*. Click **OK** in each dialog box.

Session window output

ANOVA: SolveTime versus Engineer, ProbType, Calculator

```
Factor      Type    Levels  Values
Engineer    random       6  Adams, Dixon, Erickson, Jones, Maynes, Williams
ProbType    fixed        2  Eng, Stat
Calculator  fixed        2  New, Old


Analysis of Variance for SolveTime

Source               DF      SS      MS        F      P
Engineer              5   1.053   0.211     3.13  0.039
ProbType              1  16.667  16.667   247.52  0.000
Calculator            1  72.107  72.107  1070.89  0.000
ProbType*Calculator   1   3.682   3.682    54.68  0.000
Error                15   1.010   0.067
Total                23  94.518


S = 0.259487   R-Sq = 98.93%   R-Sq(adj) = 98.36%


Means

ProbType   N  SolveTime
Eng       12     3.8250
Stat      12     5.4917


Calculator   N  SolveTime
New         12     2.9250
Old         12     6.3917


ProbType  Calculator  N  SolveTime
Eng       New         6     2.4833
Eng       Old         6     5.1667
```



```
Stat      New        6     3.3667
Stat      Old        6     7.6167
```

Interpreting the results

Minitab displays a list of factors, with their type (fixed or random), number of levels, and values. Next displayed is the analysis of variance table. The analysis of variance indicates that there is a significant calculator by problem type interaction, which implies that the decrease in mean compilation time in switching from the old to the new calculator depends upon the problem type.

Because you requested means for all factors and their combinations, the means of each factor level and factor level combinations are also displayed. These show that the mean compilation time decreased in switching from the old to new calculator type.

Example of a Mixed Model ANOVA

A company ran an experiment to see how several conditions affect the thickness of a coating substance that it manufactures. The experiment was run at two different times, in the morning and in the afternoon. Three operators were chosen from a large pool of operators employed by the company. The manufacturing process was run at three settings, 35, 44, and 52. Two determinations of thickness were made by each operator at each time and setting. Thus, the three factors are crossed. One factor, operator, is random; the other two, time and setting, are fixed.

The statistical model is:

$Y_{ijkl} = \mu + T_i + O_j + S_k + TO_{ij} + TS_{ik} + OS_{jk} + TOS_{ijk} + e_{ijkl}$,

where T_i is the time effect, O_j is the operator effect, and S_k is the setting effect, and TO_{ij}, TS_{ik}, OS_{jk}, and TOS_{ijk} are the interaction effects.

Operator, all interactions with operator, and error are random. The random terms are:

O_j TO_{ij} OS_{jk} TOS_{ijk} e_{ijkl}

These terms are all assumed to be normally distributed random variables with mean zero and variances given by

$var(O_j) = V(O)$	$var(TO_{ij}) = V(TO)$
$var(TOS_{jkl}) = V(TOS)$	$var(e_{ijkl}) = V(e) = \sigma$**2

These variances are called variance components. The output from expected means squares contains estimates of these variances.

In the unrestricted model, all these random variables are independent. The remaining terms in this model are fixed.

In the restricted model, any term which contains one or more subscripts corresponding to fixed factors is required to sum to zero over each fixed subscript. In the example, this means:

$$\sum_i (T_i) = 0 \qquad \sum_k (S_k) = 0 \qquad \sum_i (TO_{ij}) = 0$$

$$\sum_k (TS_{ik}) = 0 \qquad \sum_k (OS_{jk}) = 0 \qquad \sum_i (TOS_{ijk}) = 0$$

Your choice of model does not affect the sums of squares, degrees of freedom, mean squares, or marginal and cell means. It does affect the expected mean squares, error term for the F-tests, and the estimated variance components.

Step 1: Fit the restricted form of the model

1 Open the worksheet EXH_AOV.MTW.

2 Choose **Stat > ANOVA > Balanced ANOVA**.

3 In **Responses**, enter *Thickness*.

4 In **Model**, type *Time* | *Operator* | *Setting*.

5 In **Random Factors**, enter *Operator*.

6 Click **Options**. Check **Use the restricted form of the mixed model**. Click **OK**.

7 Click **Results**. Check **Display expected mean squares and variance components**.

8 Click **OK** in each dialog box.

Step 2: Fit the unrestricted form of the model

1 Repeat steps 1-8 above except that in 6, uncheck **Use the restricted form of the mixed model**.

Session window output for restricted case

ANOVA: Thickness versus Time, Operator, Setting

```
Factor    Type    Levels  Values
Time      fixed        2  1, 2
Operator  random       3  1, 2, 3
Setting   fixed        3  35, 44, 52


Analysis of Variance for Thickness

Source                 DF       SS      MS       F      P
Time                    1      9.0     9.0    0.29  0.644
Operator                2   1120.9   560.4  165.38  0.000
Setting                 2  15676.4  7838.2   73.18  0.001
Time*Operator           2     62.0    31.0    9.15  0.002
Time*Setting            2    114.5    57.3    2.39  0.208
Operator*Setting        4    428.4   107.1   31.61  0.000
Time*Operator*Setting   4     96.0    24.0    7.08  0.001
Error                  18     61.0     3.4
Total                  35  17568.2


S = 1.84089   R-Sq = 99.65%   R-Sq(adj) = 99.32%


                                            Expected Mean Square
                          Variance  Error  for Each Term (using
   Source                 component   term  restricted model)
1  Time                                  4  (8) + 6 (4) + 18 Q[1]
2  Operator                  46.421      8  (8) + 12 (2)
3  Setting                               6  (8) + 4 (6) + 12 Q[3]
4  Time*Operator              4.602      8  (8) + 6 (4)
5  Time*Setting                          7  (8) + 2 (7) + 6 Q[5]
6  Operator*Setting          25.931      8  (8) + 4 (6)
7  Time*Operator*Setting     10.306      8  (8) + 2 (7)
8  Error                      3.389         (8)
```

Session window output for unrestricted case

ANOVA: Thickness versus Time, Operator, Setting

```
Factor    Type    Levels  Values
Time      fixed        2  1, 2
Operator  random       3  1, 2, 3
Setting   fixed        3  35, 44, 52


Analysis of Variance for Thickness

Source                 DF       SS      MS      F      P
Time                    1      9.0     9.0   0.29  0.644
Operator                2   1120.9   560.4   4.91  0.090 x
Setting                 2  15676.4  7838.2  73.18  0.001
Time*Operator           2     62.0    31.0   1.29  0.369
Time*Setting            2    114.5    57.3   2.39  0.208
Operator*Setting        4    428.4   107.1   4.46  0.088
Time*Operator*Setting   4     96.0    24.0   7.08  0.001
Error                  18     61.0     3.4
Total                  35  17568.2

x Not an exact F-test.


S = 1.84089   R-Sq = 99.65%   R-Sq(adj) = 99.32%
```

```
                            Variance  Error  Expected Mean Square for Each
   Source                  component   term  Term (using unrestricted model)
1  Time                                   4  (8) + 2 (7) + 6 (4) + Q[1,5]
2  Operator                   37.194      *  (8) + 2 (7) + 4 (6) + 6 (4) + 12
                                             (2)
3  Setting                                6  (8) + 2 (7) + 4 (6) + Q[3,5]
4  Time*Operator               1.167      7  (8) + 2 (7) + 6 (4)
5  Time*Setting                           7  (8) + 2 (7) + Q[5]
6  Operator*Setting           20.778      7  (8) + 2 (7) + 4 (6)
7  Time*Operator*Setting      10.306      8  (8) + 2 (7)
8  Error                       3.389         (8)


* Synthesized Test.


Error Terms for Synthesized Tests

                                   Synthesis of
Source      Error DF  Error MS  Error MS
2 Operator      3.73     114.1  (4) + (6) - (7)
```

Interpreting the results

The organization of the output is the same for restricted and unrestricted models: a table of factor levels, the analysis of variance table, and as requested, the expected mean squares. The differences in the output are in the expected means squares and the F-tests for some model terms. In this example, the F-test for Operator is synthesized for the unrestricted model because it could not be calculated exactly.

Examine the 3 factor interaction, Time∗Operator∗Setting. The F-test is the same for both forms of the mixed model, giving a p-value of 0.001. This implies that the coating thickness depends upon the combination of time, operator, and setting. Many analysts would go no further than this test. If an interaction is significant, any lower order interactions and main effects involving terms of the significant interaction are not considered meaningful.

Let's examine where these models give different output. The Operator∗Setting F-test is different, because the error terms are Error in the restricted case and Time∗Operator∗Setting in the unrestricted case, giving p-values of < 0.0005 and 0.088, respectively. Likewise, the Time∗Operator differs for the same reason, giving p-values of 0.002 and 0.369, respectively, for the restricted and unrestricted cases, respectively. The estimated variance components for Operator, Time∗Operator, and Operator∗Setting also differ.

Example of a Repeated Measures Design

The following example contains data from Winer [28], p. 546, to illustrate a complex repeated measures model. An experiment was run to see how several factors affect subject accuracy in adjusting dials. Three subjects perform tests conducted at one of two noise levels. At each of three time periods, the subjects monitored three different dials and make adjustments as needed. The response is an accuracy score. The noise, time, and dial factors are crossed, fixed factors. Subject is a random factor, nested within noise. Noise is a between-subjects factor, time and dial are within-subjects factors.

We enter the model terms in a certain order so that the error terms used for the fixed factors are just below the terms for whose effects they test. (With a single random factor, the interaction of a fixed factor with the random factor becomes the error term for that fixed effect.) Because we specified Subject as Subject(Noise) the first time, we don't need to repeat "(Noise)" in the interactions involving Subject. The interaction ETime*Dial*Subject is not entered in the model because there would be zero degrees of freedom left over for error. This is the correct error term for testing ETime*Dial and by not entering ETime*Dial*Subject in the model, it is labeled as Error and we then have the error term that is needed.

1 Open the worksheet EXH_AOV.MTW.
2 Choose **Stat > ANOVA > Balanced ANOVA**.
3 In **Responses**, enter *Score*.
4 In **Model**, enter *Noise Subject(Noise) ETime Noise*ETime ETime*Subject Dial Noise*Dial Dial*Subject ETime*Dial Noise*ETime*Dial*.
5 In **Random Factors (optional)**, enter *Subject*.
6 Click **Options**.
7 Check **Use the restricted form of the mixed model**, then click **OK**.
8 Click **Results**.
9 Check **Display expected mean squares and variance components**. Click **OK** in each dialog box.

Session window output

ANOVA: Score versus Noise, ETime, Dial, Subject

```
Factor          Type    Levels  Values
Noise           fixed        2  1, 2
Subject(Noise)  random       3  1, 2, 3
ETime           fixed        3  1, 2, 3
Dial            fixed        3  1, 2, 3


Analysis of Variance for Score

Source                DF       SS       MS      F      P
Noise                  1   468.17   468.17   0.75  0.435
Subject(Noise)         4  2491.11   622.78  78.39  0.000
ETime                  2  3722.33  1861.17  63.39  0.000
Noise*ETime            2   333.00   166.50   5.67  0.029
ETime*Subject(Noise)   8   234.89    29.36   3.70  0.013
Dial                   2  2370.33  1185.17  89.82  0.000
Noise*Dial             2    50.33    25.17   1.91  0.210
Dial*Subject(Noise)    8   105.56    13.19   1.66  0.184
ETime*Dial             4    10.67     2.67   0.34  0.850
Noise*ETime*Dial       4    11.33     2.83   0.36  0.836
Error                 16   127.11     7.94
Total                 53  9924.83


S = 2.81859   R-Sq = 98.72%   R-Sq(adj) = 95.76%


                                           Expected Mean Square
                         Variance   Error  for Each Term (using
    Source               component   term  restricted model)
 1  Noise                               2  (11) + 9 (2) + 27 Q[1]
 2  Subject(Noise)          68.315     11  (11) + 9 (2)
 3  ETime                               5  (11) + 3 (5) + 18 Q[3]
 4  Noise*ETime                         5  (11) + 3 (5) + 9 Q[4]
 5  ETime*Subject(Noise)     7.139     11  (11) + 3 (5)
 6  Dial                                8  (11) + 3 (8) + 18 Q[6]
 7  Noise*Dial                          8  (11) + 3 (8) + 9 Q[7]
 8  Dial*Subject(Noise)      1.750     11  (11) + 3 (8)
 9  ETime*Dial                         11  (11) + 6 Q[9]
10  Noise*ETime*Dial                   11  (11) + 3 Q[10]
11  Error                    7.944         (11)
```

Interpreting the results

Minitab displays the table of factor levels, the analysis of variance table, and the expected mean squares. Important information to gain from the expected means squares are the estimated variance components and discovering which error term is used for testing the different model terms.

The term labeled Error is in row 11 of the expected mean squares table. The column labeled "Error term" indicates that term 11 was used to test terms 2, 5, and 8 to 10. Dial*Subject is numbered 8 and was used to test the sixth and seventh terms. You can follow the pattern for other terms.

You can gain some idea about how the design affected the sensitivity of F-tests by viewing the variance components. The variance components used in testing within-subjects factors are smaller (7.139, 1.750, 7.944) than the between-subjects variance (68.315). It is typical that a repeated measures model can detect smaller differences in means within subjects as compared to between subjects.

Of the four interactions among fixed factors, the noise by time interaction was the only one with a low p-value (0.029). This implies that there is significant evidence for judging that a subjects' sensitivity to noise changed over time. Because this interaction is significant, at least at a = 0.05, the noise and time main effects are not examined. There is also significant evidence for a dial effect (p-value < 0.0005). Among random terms, there is significant evidence for time by subject (p-value = 0.013) and subject (p-value < 0.0005) effects.

General Linear Model

Overview of Balanced ANOVA and GLM

Balanced ANOVA and general linear model (GLM) are ANOVA procedures for analyzing data collected with many different experimental designs. Your choice between these procedures depends upon the experimental design and the available options. The experimental design refers to the selection of units or subjects to measure, the assignment of treatments to these units or subjects, and the sequence of measurements taken on the units or subjects. Both procedures can fit univariate models to balanced data with up to 31 factors. Here are some of the other options:

	Balanced ANOVA	**GLM**
Can fit unbalanced data	no	yes
Can specify factors as random and obtain expected means squares	yes	yes
Fits covariates	no	yes
Performs multiple comparisons	no	yes
Fits restricted/unrestricted forms of mixed model	yes	unrestricted only

You can use balanced ANOVA to analyze data from balanced designs. See Balanced designs. You can use GLM to analyze data from any balanced design, though you cannot choose to fit the restricted case of the mixed model, which only balanced ANOVA can fit. See Restricted and unrestricted form of mixed models.

To classify your variables, determine if your factors are:

- crossed or nested
- fixed or random
- covariates

For information on how to specify the model, see Specifying the model terms, Specifying terms involving covariates, Specifying reduced models, and Specifying models for some specialized designs.

For easy entering of repeated factor levels into your worksheet, see Using patterned data to set up factor levels.

General Linear Model

Stat > ANOVA > General Linear Model

Use General Linear Model (GLM) to perform univariate analysis of variance with balanced and unbalanced designs, analysis of covariance, and regression, for each response variable.

Calculations are done using a regression approach. A "full rank" design matrix is formed from the factors and covariates and each response variable is regressed on the columns of the design matrix.

You must specify a hierarchical model. In a hierarchical model, if an interaction term is included, all lower order interactions and main effects that comprise the interaction term must appear in the model.

Factors may be crossed or nested, fixed or random. Covariates may be crossed with each other or with factors, or nested within factors. You can analyze up to 50 response variables with up to 31 factors and 50 covariates at one time. For more information see Overview of Balanced ANOVA and GLM.

Dialog box items

Responses: Select the column(s) containing the response variable(s).

Model: Specify the terms to be included in the model. See Specifying a Model for more information.

Random factors: Specify any columns containing random factors. Do not include model terms that involve other factors.

<Covariates>

<Options>

<Comparisons>

<Graphs>

<Results>

<Storage>

<Factor Plots>

Data – General Linear Model

Set up your worksheet in the same manner as with balanced ANOVA: one column for each response variable, one column for each factor, and one column for each covariate, so that there is one row for each observation. The factor columns may be numeric, text, or date/time. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. See Ordering Text Categories.

Although models can be unbalanced in GLM, they must be "full rank," that is, there must be enough data to estimate all the terms in your model. For example, suppose you have a two-factor crossed model with one empty cell. Then you can fit the model with terms A B, but not A B A∗B. Minitab will tell you if your model is not full rank. In most cases, eliminating some of the high order interactions in your model (assuming, of course, they are not important) can solve this problem.

Nesting does not need to be balanced. A nested factor must have at least 2 levels at some level of the nesting factor. If factor B is nested within factor A, there can be unequal levels of B within each level of A. In addition, the subscripts used to identify the B levels can differ within each level of A. This means, for example, that the B levels can be (1 2 3 4) in level 1 of A, (5 6 7 8) in level 2 of A, and (9 10 11 12) in level 3 of A. A nested factor must have at least 2 levels at some level of the nested factor.

If any response, factor, or covariate column contains missing data, that entire observation (row) is excluded from all computations. If you want to eliminate missing rows separately for each response, perform GLM separately for each response.

To perform an analysis using general linear model

1 Choose **Stat > ANOVA > General Linear Model**.

2 In **Responses**, enter up to 50 numeric columns containing the response variables.

3 In **Model**, type the model terms you want to fit. See Specifying the model terms.

4 If you like, use any dialog box options, then click **OK**.

Design matrix used by General Linear Model

General Linear Model uses a regression approach to fit the model that you specify. First Minitab creates a design matrix, from the factors and covariates, and the model that you specify. The columns of this matrix are the predictors for the regression.

The design matrix has n rows, where n = number of observations, and one block of columns, often called dummy variables, for each term in the model. There are as many columns in a block as there are degrees of freedom for the term. The first block is for the constant and contains just one column, a column of all ones. The block for a covariate also contains just one column, the covariate column itself.

Suppose A is a factor with 4 levels. Then it has 3 degrees of freedom and its block contains 3 columns, call them A1, A2, A3. Each row is coded as one of the following:

level of A	A1	A2	A3
1	1	0	0
2	0	1	0
3	0	0	1
4	–1	–1	–1

Suppose factor B has 3 levels nested within each level of A. Then its block contains (3 – 1) x 4 = 8 columns, call them B11, B12, B21, B22, B31, B32, B41, B42, coded as follows:

level of A	level of B	B11	B12	B21	B22	B31	B32	B41	B42
1	1	1	0	0	0	0	0	0	0
1	2	0	1	0	0	0	0	0	0
1	3	–1	–1	0	0	0	0	0	0
2	1	0	0	1	0	0	0	0	0
2	2	0	0	0	1	0	0	0	0
2	3	0	0	–1	–1	0	0	0	0
3	1	0	0	0	0	1	0	0	0
3	2	0	0	0	0	0	1	0	0
3	3	0	0	0	0	–1	–1	0	0
4	1	0	0	0	0	0	0	1	0
4	2	0	0	0	0	0	0	0	1
4	3	0	0	0	0	0	0	–1	–1

To calculate the dummy variables for an interaction term, just multiply all the corresponding dummy variables for the factors and/or covariates in the interaction. For example, suppose factor A has 6 levels, C has 3 levels, D has 4 levels, and Z and W are covariates. Then the term A ∗ C ∗ D ∗ Z ∗ W ∗ W has 5 x 2 x 3 x 1 x 1 x 1 = 30 dummy variables. To obtain them, multiply each dummy variable for A by each for C, by each for D, by the covariates Z once and W twice.

Specifying a model

Specify a model in the Model text box using the form **Y = expression**. The **Y** is not included in the Model text box. The Calc > Make Patterned Data > Simple Set of Numbers command can be helpful in entering the level numbers of a factor.

Rules for Expression Models

1 ∗ indicates an interaction term. For example, A∗B is the interaction of the factors A and B.

2 ( ) indicate nesting. When B is nested within A, type B(A). When C is nested within both A and B, type C(A B). Terms in parentheses are always factors in the model and are listed with blanks between them.

3 Abbreviate a model using a | or ! to indicate crossed factors and a − to remove terms.

Models with many terms take a long time to compute.

Examples of what to type in the Model text box

Two factors crossed: `A B A*B`

Three factors crossed: `A B C A*B A*C B*C A*B*C`

Three factors nested: `A B(A) C(A B)`

Crossed and nested (B nested within A, and both crossed with C): `A B(A) C A*C B*C(A)`

When a term contains both crossing and nesting, put the ∗ (or crossed factor) first, as in C∗B(A), not B(A)∗C

Example of entering level numbers for a data set

Here is an easy way to enter the level numbers for a three-way crossed design with a, b, and c levels of factors A, B, C, with n observations per cell:

1 Choose **Calc > Make Patterned Data > Simple Set of Numbers**, and press <F3> to reset defaults. Enter *A* in **Store patterned data in**. Enter *1* in **From first value**. Enter the number of levels in A in **To last value**. Enter the product of bcn in **List the whole sequence**. Click **OK**.

2 Choose **Calc > Make Patterned Data > Simple Set of Numbers**, and press <F3> to reset defaults. Enter *B* in **Store patterned data in**. Enter *1* in **From first value**. Enter the number of levels in B in **To last value**. Enter the number of levels in A in **List each value**. Enter the product of cn in **List the whole sequence**. Click **OK**.

3 Choose **Calc > Make Patterned Data > Simple Set of Numbers**, and press <F3> to reset defaults. Enter *C* in **Store patterned data in**. Enter *1* in *From first value*. Enter the number of levels in C in *To last value*. Enter the product of ab in **List each value**. Enter the sample size n in **List the whole sequence**. Click **OK**.

Specifying the model terms

You must specify the model terms in the **Model** box. This is an abbreviated form of the statistical model that you may see in textbooks. Because you enter the response variables in **Responses,** in **Model** you enter only the variables or products of variables that correspond to terms in the statistical model. Minitab uses a simplified version of a statistical model as it appears in many textbooks. Here are some examples of statistical models and the terms to enter in **Model**. A, B, and C represent factors.

Case	Statistical model	Terms in model
Factors A, B crossed	$y_{ijk} = \mu + a_i + b_j + ab_{ij} + e_{k(ij)}$	A B A∗ B
Factors A, B, C crossed	$y_{ijkl} = \mu + a_i + b_j + c_k + ab_{ij} + ac_{ik} + bc_{jk} + abc_{ijk} + e_{l(ijk)}$	A B C A∗B A∗ C B∗C A∗ B∗C
3 factors nested (B within A, C within A and B)	$y_{ijkl} = \mu + a_i + b_{j(i)} + c_{k(ij)} + e_{l(ijk)}$	A B(A) C(AB)
Crossed and nested (B nested within A, both crossed with C)	$y_{ijkl} = \mu + a_i + b_{j(i)} + c_k + ac_{ik} + bc_{jk(i)} + e_{l(ijk)}$	A B (A) C A∗C B∗C

In Minitab's models you omit the subscripts, μ, e, and +'s that appear in textbook models. An ∗ is used for an interaction term and parentheses are used for nesting. For example, when B is nested within A, you enter B (A), and when C is nested within both A and B, you enter C (A B). Enter B(A) C(B) for the case of 3 sequentially nested factors. Terms in parentheses are always factors in the model and are listed with blanks between them. Thus, D ∗ F (A B E) is correct but D ∗ F (A ∗ B E) and D (A ∗ B ∗ C) are not. Also, one set of parentheses cannot be used inside another set. Thus, C (A B) is correct but C (A B (A)) is not. An interaction term between a nested factor and the factor it is nested within is invalid.

See Specifying terms involving covariates for details on specifying models with covariates.

Several special rules apply to naming columns. You may omit the quotes around variable names. Because of this, variable names must start with a letter and contain only letters and numbers. Alternatively, you can use C notation (C1, C2, etc.) to denote data columns. You can use special symbols in a variable name, but then you must enclose the name in single quotes.

You can specify multiple responses. In this case, a separate analysis of variance will be performed for each response.

Specifying models for some specialized designs

Some experimental designs can effectively provide information when measurements are difficult or expensive to make or can minimize the effect of unwanted variability on treatment inference. The following is a brief discussion of three commonly used designs that will show you how to specify the model terms in Minitab. To illustrate these designs, two treatment factors (A and B) and their interaction (A∗B) are considered. These designs are not restricted to two factors, however. If your design is balanced, you can use balanced ANOVA to analyze your data. Otherwise, use GLM.

Randomized block design

A *randomized block* design is a commonly used design for minimizing the effect of variability when it is associated with discrete units (e.g. location, operator, plant, batch, time). The usual case is to randomize one replication of each treatment combination within each block. There is usually no intrinsic interest in the blocks and these are considered to be random factors. The usual assumption is that the block by treatment interaction is zero and this interaction becomes the error term for testing treatment effects. If you name the block variable as Block, enter *Block A B A∗B* in **Model** and enter *Block* in **Random Factors**.

Split-plot design

A *split-plot* design is another blocking design, which you can use if you have two or more factors. You might use this design when it is more difficult to randomize one of the factors compared to the other(s). For example, in an agricultural experiment with the factors variety and harvest date, it may be easier to plant each variety in contiguous rows and to randomly assign the harvest dates to smaller sections of the rows. The block, which can be replicated, is termed the *main plot* and within these the smaller plots (variety strips in example) are called *subplots*.

This design is frequently used in industry when it is difficult to randomize the settings on machines. For example, suppose that factors are temperature and material amount, but it is difficult to change the temperature setting. If the blocking factor is operator, observations will be made at different temperatures with each operator, but the temperature setting is held constant until the experiment is run for all material amounts. In this example, the plots under operator constitute the main plots and temperatures constitute the subplots.

There is no single error term for testing all factor effects in a split-plot design. If the levels of factor A form the subplots, then the mean square for Block ∗ A will be the error term for testing factor A. There are two schools of thought for what should be the error term to use for testing B and A ∗ B. If you enter the term Block ∗ B, the expected mean squares show that the mean square for Block ∗ B is the proper term for testing factor B and that the remaining error (which is Block ∗ A ∗ B) will be used for testing A ∗ B. However, it is often assumed that the Block ∗ B and Block ∗ A ∗ B interactions do not exist and these are then lumped together into error [6]. You might also pool the two terms if the mean square for Block ∗ B is small relative to Block ∗ A ∗ B. If you don't pool, enter *Block A Block ∗ A B Block ∗B A ∗ B* in **Model** and what is labeled as Error is really Block ∗ A ∗ B. If you do pool terms, enter *Block A Block ∗ A B A ∗ B* in **Model** and what is labeled as Error is the set of pooled terms. In both cases enter *Block* in **Random Factors**.

Latin square with repeated measures design

A *repeated measures* design is a design where repeated measurements are made on the same subject. There are a number of ways in which treatments can be assigned to subjects. With living subjects especially, systematic differences (due to learning, acclimation, resistance, etc.) between successive observations may be suspected. One common way to assign treatments to subjects is to use a Latin square design. An advantage of this design for a repeated measures experiment is that it ensures a balanced fraction of a complete factorial (i.e. all treatment combinations represented) when subjects are limited and the sequence effect of treatment can be considered to be negligible.

A *Latin square* design is a blocking design with two orthogonal blocking variables. In an agricultural experiment there might be perpendicular gradients that might lead you to choose this design. For a repeated measures experiment, one blocking variable is the group of subjects and the other is time. If the treatment factor B has three levels, b1, b2, and b3, then one of twelve possible Latin square randomizations of the levels of B to subjects groups over time is:

	Time 1	Time 2	Time 3
Group 1	b2	b3	b1
Group 2	b3	b1	b2
Group 3	b1	b2	b3

The subjects receive the treatment levels in the order specified across the row. In this example, group 1 subjects would receive the treatments levels in order b_2, b_3, b_1. The interval between administering treatments should be chosen to minimize carryover effect of the previous treatment.

This design is commonly modified to provide information on one or more additional factors. If each group was assigned a different level of factor A, then information on the A and A ∗ B effects could be made available with minimal effort if an assumption about the sequence effect given to the groups can be made. If the sequence effects are negligible compared to the effects of factor A, then the group effect could be attributed to factor A. If interactions with time are negligible, then partial information on the A ∗ B interaction may be obtained [29]. In the language of repeated measures designs, factor A is called a *between-subjects* factor and factor B a *within-subjects* factor.

Let's consider how to enter the model terms into Minitab. If the group or A factor, subject, and time variables were named A, Subject, and Time, respectively, enter *A Subject(A) Time B A ∗ B* in **Model** and enter *Subject* in **Random Factors**.

It is not necessary to randomize a repeated measures experiments according to a Latin square design. See Example of a repeated measures design for a repeated measures experiment where the fixed factors are arranged in a complete factorial design.

Specifying reduced models

You can fit *reduced* models. For example, suppose you have a three factor design, with factors, A, B, and C. The *full* model would include all one factor terms: A, B, C, all two-factor interactions: A ∗ B, A ∗ C, B ∗ C, and the three-factor interaction: A ∗ □B ∗ C. It becomes a reduced model by omitting terms. You might reduce a model if terms are not significant or if you need additional error degrees of freedom and you can assume that certain terms are zero. For this example, the model with terms A B C A ∗ B is a reduced three-factor model.

One rule about specifying reduced models is that they must be hierarchical. That is, for a term to be in the model, all lower order terms contained in it must also be in the model. For example, suppose there is a model with four factors: A, B, C, and D. If the term A ∗ B ∗ C is in the model then the terms A B C A ∗ B A ∗ C B ∗ C must also be in the model, though any terms with D do not have to be in the model. The hierarchical structure applies to nesting as well. If B (A) is in the model, then A must be also.

Because models can be quite long and tedious to type, two shortcuts have been provided. A vertical bar indicates crossed factors, and a minus sign removes terms.

Long form	**Short form**
A B C A ∗ B A ∗ C B ∗ C A ∗ B ∗ C	A | B | C
A B C A ∗ B A ∗ C B ∗ C	A | B | C – A ∗ B ∗ C
A B C B ∗ C E	A B | C E
A B C D A ∗ B A ∗ C A ∗ D B ∗ C B ∗ D C ∗ D A ∗ B ∗ D A ∗ C ∗ D B ∗ C ∗ D	A | B | C | D – A ∗ B ∗ C – A ∗ B ∗ C ∗ D
A B (A) C A ∗ C B ∗ C	A | B (A) | C

In general, all crossings are done for factors separated by bars unless the cross results in an illegal term. For example, in the last example, the potential term A ∗ B (A) is illegal and Minitab automatically omits it. If a factor is nested, you must indicate this when using the vertical bar, as in the last example with the term B (A).

Using patterned data to set up factor levels

Minitab's set patterned data capability can be helpful when entering numeric factor levels. For example, to enter the level values for a three-way crossed design with a, b, and c (a, b, and c represent numbers) levels of factors A, B, C, and n observations per cell, fill out the Calc > Set Patterned Data > Simple Set of Numbers dialog box and execute 3 times, once for each factor, as shown:

	Factor		
Dialog item	**A**	**B**	**C**
From first value	1	1	1
From last value	a	b	c
List each value	bcn	cn	n
List the whole sequence	1	a	ab

Coefficients in general linear models

General Linear Model (GLM) uses a regression approach to fit your model. First, GLM codes the factor levels as dummy or indicator variables using a 1, 0, – 1, coding scheme. For more information on how Minitab codes the data for a GLM

analysis, see Design matrix used by General Linear Model. The dummy variables are then used to calculate the coefficients for all terms. In GLM, the coefficients represent the distance between factor levels and the overall mean.

You can view the ANOVA model equation by displaying the table of coefficients for all terms in the GLM output. In the Results subdialog box, choose **In addition, coefficients for all terms**.

After you conduct the analysis, you will notice that coefficients are listed for all but one of the levels for each factor. This level is the reference level or baseline. All estimated coefficients are interpreted relative to the reference level. In some cases, you may want to know the reference level coefficient to understand how the reference value compares in size and direction to the overall mean.

Suppose you perform a general linear model test with 2 factors. Factor 1 has 3 levels (A, B, and C), and Factor 2 has 2 levels (High and Low). Minitab codes these levels using indicator variables. For factor 1: A = 1, B = 0, and C = – 1. For Factor 2: High = 1 and Low = – 1.

You obtain the following table of coefficients:

```
Term          Coef   SE Coef        T       P
Constant    5.0000    0.1954    25.58   0.000
Factor1
A          -3.0000    0.2764   -10.85   0.000
B          -0.5000    0.2764    -1.81   0.108
Factor2
High       -0.8333    0.1954    -4.26   0.003
```

The ANOVA model is: Response = 5.0 – 3.0 ∗ A – 0.5 ∗ B – 0.833 ∗ High

Notice that the table does not include the coefficients for C (Factor 1) or Low (Factor 2), which are the reference levels for each factor. However, you can easily calculate these values by subtracting the overall mean from each level mean. The constant term is the overall mean. Use Stat > Basic Statistics > Display Descriptive Statistics to obtain the mean for each level. The means are:

Overall	5.0
A (Factor 1)	2.0
B (Factor 1)	4.5
C (Factor 1)	8.5
High (Factor 2)	4.1667
High (Factor 2)	5.8333

The coefficients are calculated as the level mean – overall mean, Thus, the coefficients for each level are:

Level A effect = 2.0 – 5.0 = – 3.0

Level B effect = 4.5 – 5.0 = – 0.5

Level C effect = 8.5 – 5.0 = 3.5 (not given in the coefficients table)

Level High effect = 4.1667 – 5.0 = – 0.8333

Level Low effect = 5.8333 – 5.0 = 0.8333 (not given in the coefficients table)

Tip A quick way to obtain the coefficients not listed in the table is by adding all of the level coefficients for a factor (excluding the intercept) and multiplying by – 1. For example, the coefficient for Level C = – 1 * [(– 3.0) + (– 0.50)] = 3.5.

If you add a covariate or have unequal sample sizes within each group, coefficients are based on weighted means for each factor level rather than the arithmetic mean (sum of the observations divided by n).

General Linear Model – Covariates

Stat > ANOVA > General Linear Model > Covariates

Enter covariates into the model. These are entered into the model first by default.

Dialog box items

Covariates: Enter columns containing the covariates.

Specifying terms involving covariates

You can specify variables to be covariates in GLM. You must specify the covariates in *Covariates*, but you can enter the covariates in *Model*, though this is not necessary unless you cross or nest the covariates (see table below).

In an unbalanced design or a design involving covariates, GLM's sequential sums of squares (the additional model sums of squares explained by a variable) will depend upon the order in which variables enter the model. If you do not enter the covariates in *Model* when using GLM, they will be fit first, which is what you usually want when a covariate contributes background variability. The subsequent order of fitting is the order of terms in *Model*. The sequential sums of squares for unbalanced terms A B will be different depending upon the order that you enter them in the model. The default adjusted sums of squares (sums of squares with all other terms in the model), however, will be the same, regardless of model order.

GLM allows terms containing covariates crossed with each other and with factors, and covariates nested within factors. Here are some examples of these models, where A is a factor.

Case	Covariates	Terms in model
test homogeneity of slopes (covariate crossed with factor)	X	A X A∗X
same as previous	X	A | X
quadratic in covariate (covariate crossed with itself)	X	A X X∗X
full quadratic in two covariates (covariates crossed)	X Z	A X Z X∗X Z∗Z X∗Z
separate slopes for each level of A (covariate nested within a factor)	X	A X (A)

General Linear Model – Options

Stat > ANOVA > General Linear Model > Options

Allows you to choose a weighted fit and the sums of squares type used in the ANOVA.

Dialog box items

Do a weighted fit, using weights in: Enter a column of weights for a weighted fit. See Weighted regression for more information.

Sum of Squares: Select a sums of squares for calculating F-and p-values.

Adjusted (Type III): Choose if you want sums of squares for terms with other terms in the model

Sequential (Type I): Choose if you want sums of squares with only previous terms in the model

Adjusted vs. sequential sums of squares

Minitab by default uses adjusted (Type III) sums of squares for all GLM calculations. Adjusted sums of squares are the additional sums of squares determined by adding each particular term to the model given the other terms are already in the model. You also have the choice of using sequential (Type I) sums of squares in all GLM calculations. Sequential sums of squares are the sums of squares added by a term with only the previous terms entered in the model. These sums of squares can differ when your design is unbalanced or if you have covariates. Usually, you would probably use adjusted sums of squares. However, there may be cases where you might want to use sequential sums of squares.

General Linear Model – Comparisons

Stat > ANOVA > General Linear Model > Comparisons

Specify terms for comparing the means, as well as the type of multiple comparisons.

Dialog box items

Pairwise comparisons: Choose to obtain pairwise comparison of all mean for designated terms.

Comparisons with a control: Choose to obtain comparisons of means with the mean of a control level.

Terms: Enter the model terms for comparison.

Control levels: Enter the control level if you chose comparisons with a control. (IMPORTANT: For text variables, you must enclose factor levels in double quotes, even if there are no spaces in them.)

Method: Select the multiple comparison method(s). See Multiple comparisons.

Tukey: Choose the Tukey (also call Tukey-Kramer method in unbalanced case) method.

Dunnett: Choose the Dunnett method.

Bonferroni: Choose the Bonferroni method.

Sidak: Choose the Sidak method.

Alternative: Choose one of three possible alternative hypotheses when you choose comparisons with a control. The null hypothesis is equality of treatment and control means.

Less than: Choose the alternative hypothesis of the treatment mean being less than the mean of the control group.

Not equal: Choose the alternative hypothesis of the treatment mean being not equal to the mean of the control group.

Greater than: Choose the alternative hypothesis of the treatment mean being greater than the mean of the control group.

Confidence interval, with confidence level: Check to specify a confidence level and then enter a value for the intervals that is between 0 and 100 (the default is 95%).

Test: Check to select the hypothesis test form of multiple comparison output.

Multiple comparisons of means

Multiple comparisons of means allow you to examine which means are different and to estimate by how much they are different. When you have multiple factors, you can obtain multiple comparisons of means through GLM's Comparisons subdialog box.

There are some common pitfalls to the use of multiple comparisons. If you have a quantitative factor you should probably examine linear and higher order effects rather than performing multiple comparisons (see [12] and Example of using GLM to fit linear and quadratic effects). In addition, performing multiple comparisons for those factors which appear to have the greatest effect or only those with a significant F-test can result in erroneous conclusions (see Which means to compare? below).

You have the following choices when using multiple comparisons:

- Pairwise comparisons or comparisons with a control
- Which means to compare
- The method of comparison
- Display comparisons in confidence interval or hypothesis test form
- The confidence level, if you choose to display confidence intervals
- The alternative, if you choose comparisons with a control

Following are some guidelines for making these choices.

Pairwise comparisons or comparison with a control

Choose **Pairwise Comparisons** when you do not have a control level but you would like to examine which pairs of means are different.

Choose **Comparisons with a Control** when you are comparing treatments to a control. When this method is suitable, it is inefficient to use the all-pairwise approach, because the all-pairwise confidence intervals will be wider and the hypothesis tests less powerful for a given family error rate. If you do not specify a level that represents the control, Minitab will assume that the lowest level of the factors is the control. If you wish to change which level is the control, specify a level that represents the control for each term that you are comparing the means of. If these levels are text or date/time, enclose each with double quotes.

Which means to compare

Choosing which means to compare is an important consideration when using multiple comparisons; a poor choice can result in confidence levels that are not what you think. Issues that should be considered when making this choice might include: 1) should you compare the means for only those terms with a significant F-test or for those sets of means for which differences appear to be large? 2) how deep into the design should you compare means–only within each factor, within each combination of first-level interactions, or across combinations of higher level interactions?

It is probably a good idea to decide which means you will compare before collecting your data. If you compare only those means with differences that appear to be large, which is called data snooping, then you are increasing the likelihood that the results suggest a real difference where no difference exists [9], [19]. Similarly, if you condition the application of multiple comparisons upon achieving a significant F-test, then the error rate of the multiple comparisons can be higher than the error rate in the unconditioned application of multiple comparisons [9], [15]. The multiple comparison methods have protection against false positives already built in.

In practice, however, many people commonly use F-tests to guide the choice of which means to compare. The ANOVA F-tests and multiple comparisons are not entirely separate assessments. For example, if the p-value of an F-test is 0.9, you probably will not find statistically significant differences among means by multiple comparisons.

How deep within the design should you compare means? There is a trade-off: if you compare means at all two-factor combinations and higher orders turn out to be significant, then the means that you compare might be a mix of effects; if you compare means at too deep a level, you lose power because the sample sizes become smaller and the number of

comparisons become larger. In practice, you might decide to compare means for factor level combinations for which you believe the interactions are meaningful.

Minitab restricts the terms that you can compare means for to fixed terms or interactions among fixed terms. Nesting is considered to be a form of interaction.

To specify which means to compare, enter terms from the model in the **Terms** box. If you have 2 factors named A and B, entering A B will result in multiple comparisons within each factor. Entering A ∗ B will result in multiple comparisons for all level combination of factors A and B. You can use the notation A | B to indicate interaction for pairwise comparisons but not for comparisons with a control.

The multiple comparison method

You can choose from among three methods for both pairwise comparisons and comparisons with a control. Each method provides simultaneous or joint confidence intervals, meaning that the confidence level applies to the set of intervals computed by each method and not to each one individual interval. By protecting against false positives with multiple comparisons, the intervals are wider than if there were no protection.

The Tukey (also called Tukey-Kramer in the unbalanced case) and Dunnett methods are extensions of the methods used by one-way ANOVA. The Tukey approximation has been proven to be conservative when comparing three means. "Conservative" means that the true error rate is less than the stated one. In comparing larger numbers of means, there is no proof that the Tukey method is conservative for the general linear model. The Dunnett method uses a factor analytic method to approximate the probabilities of the comparisons. Because it uses the factor analytic approximation, the Dunnett method is not generally conservative. The Bonferroni and Sidak methods are conservative methods based upon probability inequalities. The Sidak method is slightly less conservative than the Bonferroni method.

Some characteristics of the multiple comparison methods are summarized below:

Comparison method	Properties
Dunnett	comparison to a control only, not proven to be conservative
Tukey	all pairwise differences only, not proven to be conservative
Bonferroni	most conservative
Sidak	conservative, but slightly less so than Bonferroni

Display of comparisons in confidence interval or hypothesis test form

Minitab presents multiple comparison results in confidence interval and/or hypothesis test form. Both are given by default.

When viewing confidence intervals, you can assess the practical significance of differences among means, in addition to statistical significance. As usual, the null hypothesis of no difference between means is rejected if and only if zero is not contained in the confidence interval. When you request confidence intervals, you can specify family confidence levels for the confidence intervals. The default level is 95%.

Minitab calculates adjusted p-values for hypothesis test statistics. The adjusted p-value for a particular hypothesis within a collection of hypotheses is the smallest family wise a level at which the particular hypothesis would be rejected.

General Linear Model – Graphs

Stat > ANOVA > General Linear Model > Graphs

Displays residual plots. You do not have to store the residuals and fits in order to produce these plots.

Dialog box items

Residuals for Plots: You can specify the type of residual to display on the residual plots.

- **Regular:** Choose to plot the regular or raw residuals.
- **Standardized:** Choose to plot the standardized residuals.
- **Deleted:** Choose to plot the Studentized deleted residuals.

Residual Plots

- **Individual plots:** Choose to display one or more plots.
  - **Histogram of residuals:** Check to display a histogram of the residuals.
  - **Normal plot of residuals:** Check to display a normal probability plot of the residuals.
  - **Residuals versus fits:** Check to plot the residuals versus the fitted values.
  - **Residuals versus order:** Check to plot the residuals versus the order of the data. The row number for each data point is shown on the x-axis–for example, 1 2 3 4... n.

Four in one: Choose to display a layout of a histogram of residuals, a normal plot of residuals, a plot of residuals versus fits, and a plot of residuals versus order.

Residuals versus the variables: Enter one or more columns containing the variables against which you want to plot the residuals. Minitab displays a separate graph for each column.

General Linear Model – Results

Stat > ANOVA > General Linear Model > Results

Control the display of results to the Session window, the display of expected means squares and variance components, and display means of model term levels.

Dialog box items

Display of Results

Display nothing: Choose to display nothing.

Analysis of variance table: Choose to display only the analysis of variance table.

In addition, coefficient for covariate terms and table of unusual observations: Choose to display, in addition to the ANOVA table, a table of covariate term coefficients and a table of unusual observations.

In addition, coefficient for all terms: Choose to display, in addition to the above tables, the coefficients for all terms.

Display expected mean squares and variance components: Check to display the expected mean squares and variance component estimates for random terms.

Display means corresponding to the terms: Enter the model terms for which to display least squares means and their standard errors.

General Linear Model – Storage

Stat > ANOVA > General Linear Model > Storage

Stores the residuals, fitted values, and many other diagnostics for further analysis (see Checking your model).

Dialog box items

Diagnostic Measures

Residuals: Check to store the residuals.

Standardized residuals: Check to store the standardized residuals.

Deleted t residuals: Check to store Studentized residuals.

Hi [leverage]: Check to store leverages.

Cook's distance: Check to store Cook's distance.

DFITS: Check to store DFITS.

Characteristics of Estimated Equation

Coefficients: Check to store the coefficients for a model that corresponds to the design matrix. (If M1 contains the design matrix and C1 the coefficients, then M1 times C1 gives the fitted values.)

Fits: Check to store the fitted values.

Design matrix: Check to store the design matrix corresponding to your model.

General Linear Model – Factorial Plots

Stat > ANOVA > General Linear Model > Factor Plots

Displays plots of the main effects and interactions in your data.

Dialog box items

Main Effects Plot: Display plots of main effects.

Factors: Chose the factors to plot.

Minimum for Y (response) scale: Replace the default minimum value for the Y-axis with one you chose.

Maximum for Y (response) scale: Replace the default maximum value for the Y-axis with one you chose.

Title: Replace the default title with one of your own.

Interactions Plot: Display plots of two-way interactions.

Factors: Choose the factors to include in the plot(s).

Display full interaction plot matrix: By default, Minitab chooses one factor to represent on the X-axis and represents the levels of the other factor with different symbols and lines. Check this option to display each interaction twice, once with each factor represented on the X-axis.

Minimum for Y (response) scale: Replace the default minimum value for the Y-axis with one you chose.

Maximum for Y (response) scale: Replace the default maximum value for the Y-axis with one you chose.

Title: Replace the default title with one of your own.

Example of Using GLM to fit Linear and Quadratic Effects

An experiment is conducted to test the effect of temperature and glass type upon the light output of an oscilloscope. There are three glass types and three temperature levels: 100, 125, and 150 degrees Fahrenheit. These factors are fixed because we are interested in examining the response at those levels. The example and data are from Montgomery [14], page 252.

When a factor is quantitative with three or more levels it is appropriate to partition the sums of squares from that factor into effects of polynomial orders [11]. If there are k levels to the factor, you can partition the sums of squares into k-1 polynomial orders. In this example, the effect due to the quantitative variable temperature can be partitioned into linear and quadratic effects. Similarly, you can partition the interaction. To do this, you must code the quantitative variable with the actual treatment values (that is, code Temperature levels as 100, 125, and 150), use GLM to analyze your data, and declare the quantitative variable to be a covariate.

1 Open the worksheet EXH_AOV.MTW.
2 Choose **Stat > ANOVA > General Linear Model**.
3 In **Responses**, enter *LightOutput*.
4 In **Model**, type *Temperature Temperature ∗Temperature GlassType GlassType ∗Temperature GlassType ∗ Temperature ∗Temperature*.
5 Click **Covariates**. In **Covariates**, enter *Temperature*.
6 Click **OK** in each dialog box.

Session window output

General Linear Model: LightOutput versus GlassType

```
Factor     Type   Levels  Values
GlassType  fixed       3  1, 2, 3


Analysis of Variance for LightOutput, using Adjusted SS for Tests

Source                            DF   Seq SS  Adj SS  Adj MS       F      P
Temperature                        1  1779756  262884  262884  719.21  0.000
Temperature*Temperature            1   190579  190579  190579  521.39  0.000
GlassType                          2   150865   41416   20708   56.65  0.000
GlassType*Temperature              2   226178   51126   25563   69.94  0.000
GlassType*Temperature*Temperature  2    64374   64374   32187   88.06  0.000
Error                             18     6579    6579     366
Total                             26  2418330


S = 19.1185   R-Sq = 99.73%   R-Sq(adj) = 99.61%


Term                                  Coef  SE Coef       T      P
Constant                           -4968.8    191.3  -25.97  0.000
Temperature                         83.867    3.127   26.82  0.000
Temperature*Temperature           -0.28516  0.01249  -22.83  0.000
Temperature*GlassType
            1                      -24.400    4.423   -5.52  0.000
            2                      -27.867    4.423   -6.30  0.000
Temperature*Temperature*GlassType
                        1          0.11236  0.01766    6.36  0.000
                        2          0.12196  0.01766    6.91  0.000


Unusual Observations for LightOutput
```



```
Obs  LightOutput      Fit  SE Fit  Residual  St Resid
 11      1070.00  1035.00   11.04     35.00      2.24 R
 17      1000.00  1035.00   11.04    -35.00     -2.24 R

R denotes an observation with a large standardized residual.
```

Interpreting the results

Minitab first displays a table of factors, with their number of levels, and the level values. The second table gives an analysis of variance table. This is followed by a table of coefficients, and then a table of unusual observations.

The Analysis of Variance table gives, for each term in the model, the degrees of freedom, the sequential sums of squares (Seq SS), the adjusted (partial) sums of squares (Adj SS), the adjusted means squares (Adj MS), the F-statistic from the adjusted means squares, and its p-value. The sequential sums of squares is the added sums of squares given that prior terms are in the model. These values depend upon the model order. The adjusted sums of squares are the sums of squares given that all other terms are in the model. These values do not depend upon the model order. If you had selected sequential sums of squares in the Options subdialog box, Minitab would use these values for mean squares and F-tests.

In the example, all p-values were printed as 0.000, meaning that they are less than 0.0005. This indicates significant evidence of effects if your level of significance, a, is greater than 0.0005. The significant interaction effects of glass type with both linear and quadratic temperature terms implies that the coefficients of second order regression models of the effect of temperature upon light output depends upon the glass type.

The next table gives the estimated coefficients for the covariate, Temperature, and the interactions of Temperature with GlassType, their standard errors, t-statistics, and p-values. Following the table of coefficients is a table of unusual values. Observations with large standardized residuals or large leverage values are flagged. In our example, two values have standardized residuals whose absolute values are greater than 2.

Example of Using GLM and Multiple Comparisons with an Unbalanced Nested Design

Four chemical companies produce insecticides that can be used to kill mosquitoes, but the composition of the insecticides differs from company to company. An experiment is conducted to test the efficacy of the insecticides by placing 400 mosquitoes inside a glass container treated with a single insecticide and counting the live mosquitoes 4 hours later. Three replications are performed for each product. The goal is to compare the product effectiveness of the different companies. The factors are fixed because you are interested in comparing the particular brands. The factors are nested because each insecticide for each company is unique. The example and data are from Milliken and Johnson [12], page 414. You use GLM to analyze your data because the design is unbalanced and you will use multiple comparisons to compare the mean response for the company brands.

1 Open the worksheet EXH_AOV.MTW.
2 Choose **Stat > ANOVA > General Linear Model**.
3 In **Responses**, enter *NMosquito*.
4 In **Model**, type *Company Product(Company)*.
5 Click **Comparisons**. Under **Pairwise Comparisons**, enter *Company* in **Terms**.
6 Under **Method**, check **Tukey**. Click **OK** in each dialog box.

Session window output

General Linear Model: NMosquito versus Company, Product

```
Factor            Type   Levels  Values
Company           fixed       4  A, B, C, D
Product(Company)  fixed      11  A1, A2, A3, B1, B2, C1, C2, D1, D2, D3, D4


Analysis of Variance for NMosquito, using Adjusted SS for Tests

Source            DF   Seq SS   Adj SS  Adj MS       F      P
Company            3  22813.3  22813.3  7604.4  132.78  0.000
Product(Company)   7   1500.6   1500.6   214.4    3.74  0.008
Error             22   1260.0   1260.0    57.3
Total             32  25573.9


S = 7.56787   R-Sq = 95.07%   R-Sq(adj) = 92.83%
```

```
Tukey 95.0% Simultaneous Confidence Intervals
Response Variable NMosquito
All Pairwise Comparisons among Levels of Company
Company = A  subtracted from:

Company   Lower  Center   Upper  --------+---------+---------+--------
B         -2.92    8.17   19.25                              (---*----)
C        -52.25  -41.17  -30.08         (----*---)
D        -61.69  -52.42  -43.14     (---*---)
                                 --------+---------+---------+--------
                                       -50       -25         0


Company = B  subtracted from:

Company   Lower  Center   Upper  --------+---------+---------+--------
C        -61.48  -49.33  -37.19     (----*----)
D        -71.10  -60.58  -50.07  (---*---)
                                 --------+---------+---------+--------
                                       -50       -25         0


Company = C  subtracted from:

Company   Lower  Center    Upper  --------+---------+---------+--------
D        -21.77  -11.25  -0.7347                     (----*---)
                                  --------+---------+---------+--------
                                        -50       -25         0


Tukey Simultaneous Tests
Response Variable NMosquito
All Pairwise Comparisons among Levels of Company
Company = A  subtracted from:

         Difference       SE of           Adjusted
Company    of Means  Difference  T-Value   P-Value
B              8.17       3.989     2.05    0.2016
C            -41.17       3.989   -10.32    0.0000
D            -52.42       3.337   -15.71    0.0000


Company = B  subtracted from:

         Difference       SE of           Adjusted
Company    of Means  Difference  T-Value   P-Value
C            -49.33       4.369   -11.29    0.0000
D            -60.58       3.784   -16.01    0.0000


Company = C  subtracted from:

         Difference       SE of           Adjusted
Company    of Means  Difference  T-Value   P-Value
D            -11.25       3.784   -2.973    0.0329
```

Interpreting the results

Minitab displays a factor level table, an ANOVA table, multiple comparison confidence intervals for pairwise differences between companies, and the corresponding multiple comparison hypothesis tests. The ANOVA F-tests indicate that there is significant evidence for company effects.

Examine the multiple comparison confidence intervals. There are three sets: 1) for the company A mean subtracted from the company B, C, and D means; 2) for the company B mean subtracted from the company C and D means; and 3) for the company C mean subtracted from the company D mean. The first interval, for the company B mean minus the company A mean, contains zero is in the confidence interval. Thus, there is no significant evidence at $\alpha = 0.05$ for differences in means. However, there is evidence that all other pairs of means are different, because the confidence intervals for the differences in means do not contain zero. An advantage of confidence intervals is that you can see the magnitude of the differences between the means.

Examine the multiple comparison hypothesis tests. These are laid out in the same way as the confidence intervals. You can see at a glance the mean pairs for which there is significant evidence of differences. The adjusted p-values are small for all but one comparison, that of company A to company B. An advantage of hypothesis tests is that you can see what a level would be required for significant evidence of differences.

Fully Nested ANOVA

Fully Nested ANOVA

Stat > ANOVA > Fully Nested ANOVA

Use to perform fully nested (hierarchical) analysis of variance and to estimate variance components for each response variable. All factors are implicitly assumed to be random. Minitab uses sequential (Type I) sums of squares for all calculations.

You can analyze up to 50 response variables with up to 9 factors at one time.

If your design is not hierarchically nested or if you have fixed factors, use either Balanced ANOVA or GLM. Use GLM if you want to use adjusted sums of squares for a fully nested model.

Dialog box items

Responses: Enter the columns containing your response variables.

Factors: Enter the columns containing the factors in hierarchical order.

Data – Fully Nested ANOVA

Set up your worksheet in the same manner as with Balanced ANOVA or GLM: one column for each response variable and one column for each factor, so that there is one row for each observation. The factor columns may be numeric, text, or date/time. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. See Ordering Text Categories.

Nesting does not need to be balanced. A nested factor must have at least 2 levels at some level of the nesting factor. If factor B is nested within factor A, there can be unequal levels of B within each level of A. In addition, the subscripts used to identify the B levels can differ within each level of A.

If any response or factor column contains missing data, that entire observation (row) is excluded from all computations. If an observation is missing for one response variable, that row is eliminated for all responses. If you want to eliminate missing rows separately for each response, perform a fully nested ANOVA separately for each response.

You can analyze up to 50 response variables with up to 9 factors at one time.

To perform an analysis using fully nested ANOVA

1 Choose **Stat > ANOVA > Fully Nested ANOVA**.
2 In **Responses**, enter up to 50 numeric columns containing the response variables.
3 In **Factors**, type in the factors in hierarchical order. See Fully Nested or Hierarchical Models.
4 Click **OK**.

Fully Nested or Hierarchical Models

Minitab fits a fully nested or hierarchical model with the nesting performed according to the order of factors in the **Factors** box. If you enter factors A B C, then the model terms will be A B(A) C(B). You do not need to specify these terms in model form as you would for Balanced ANOVA or GLM.

Minitab uses sequential (Type I) sums of squares for all calculations of fully nested ANOVA. This usually makes sense for a hierarchical model. General Linear Models (GLM) offers the choice of sequential or adjusted (Type III) sums of squares and uses the adjusted sums of squares by default. These sums of squares can differ when your design is unbalanced. Use GLM if you want to use adjusted sums of squares for calculations.

Example of Fully Nested ANOVA

You are an engineer trying to understand the sources of variability in the manufacture of glass jars. The process of making the glass requires mixing materials in small furnaces for which the temperature setting is to be 475° F. Your company has a number of plants where the jars are made, so you select four as a random sample. You conduct an experiment and measure furnace temperature for four operators over four different shifts. You take two batch measurements during each shift. Because your design is fully nested, you use Fully Nested ANOVA to analyze your data.

1 Open the worksheet FURNTEMP.MTW.

2 Choose **Stat > ANOVA > Fully Nested ANOVA**.

3 In **Responses**, enter *Temp*.

4 In **Factors**, enter *Plant - Batch*. Click **OK**.

Session window output

Nested ANOVA: Temp versus Plant, Operator, Shift, Batch

```
Analysis of Variance for Temp

Source      DF         SS         MS      F      P
Plant        3   731.5156   243.8385  5.854  0.011
Operator    12   499.8125    41.6510  1.303  0.248
Shift       48  1534.9167    31.9774  2.578  0.000
Batch      128  1588.0000    12.4062
Total      191  4354.2448


Variance Components

                      % of
Source    Var Comp.  Total  StDev
Plant         4.212  17.59  2.052
Operator      0.806   3.37  0.898
Shift         6.524  27.24  2.554
Batch        12.406  51.80  3.522
Total        23.948         4.894


Expected Mean Squares

1  Plant       1.00(4) +  3.00(3) + 12.00(2) + 48.00(1)
2  Operator    1.00(4) +  3.00(3) + 12.00(2)
3  Shift       1.00(4) +  3.00(3)
4  Batch       1.00(4)
```

Interpreting the results

Minitab displays three tables of output: 1) the ANOVA table, 2) the estimated variance components, and 3) the expected means squares. There are four sequentially nested sources of variability in this experiment: plant, operator, shift, and batch. The ANOVA table indicates that there is significant evidence for plant and shift effects at $\alpha = 0.05$ (F-test p-values < 0.05). There is no significant evidence for an operator effect. The variance component estimates indicate that the variability attributable to batches, shifts, and plants was 52, 27, and 18 percent, respectively, of the total variability.

If a variance component estimate is less than zero, Minitab displays what the estimate is, but sets the estimate to zero in calculating the percent of total variability.

Balanced MANOVA

Balanced MANOVA

Stat > ANOVA > Balanced MANOVA

Use balanced MANOVA to perform multivariate analysis of variance (MANOVA) for balanced designs. You can take advantage of the data covariance structure to simultaneously test the equality of means from different responses.

Your design must be balanced, with the exception of one-way designs. *Balanced* means that all treatment combinations (cells) must have the same number of observations. Use General MANOVA to analyze either balanced and unbalanced MANOVA designs or if you have covariates. You cannot designate factors to be random with general MANOVA, unlike for balanced ANOVA, though you can work around this restriction by supplying error terms to test the model terms.

Factors may be crossed or nested, fixed or random.

Dialog box items

Responses: Enter up to 50 numeric columns containing the response variables

Model: Type the model terms that you want to fit.

Random Factors: Enter which factors are random factors.

<Options>

<Graphs>

<Results>

<Storage>

Data – Balanced MANOVA

You need one column for each response variable and one column for each factor, with each row representing an observation. Regardless of whether factors are crossed or nested, use the same form for the data. Factor columns may be numeric, text, or date/time. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. See Ordering Text Categories. You may include up to 50 response variables and up to 31 factors at one time.

Balanced data are required except for one-way designs. The requirement for balanced data extends to nested factors as well. Suppose A has 3 levels, and B is nested within A. If B has 4 levels within the first level of A, B must have 4 levels within the second and third levels of A. Minitab will tell you if you have unbalanced nesting. In addition, the subscripts used to indicate the 4 levels of B within each level of A must be the same. Thus, the four levels of B cannot be (1 2 3 4) in level 1 of A, (5 6 7 8) in level 2 of A, and (9 10 11 12) in level 3 of A. You can use general MANOVA if you have different levels of B within the levels of A.

If any response or factor column specified contains missing data, that entire observation (row) is excluded from all computations. The requirement that data be balanced must be preserved after missing data are omitted.

To perform a balanced MANOVA

1 Choose **Stat > ANOVA > Balanced MANOVA**.

2 In **Responses**, enter up to 50 numeric columns containing the response variables.

3 In **Model**, type the model terms that you want to fit. See Overview of Balanced ANOVA and GLM.

4 If you like, use any dialog box options, then click **OK**.

Balanced MANOVA – Options

Stat > ANOVA > Balanced MANOVA > Options

Dialog box items

Use the restricted form of the model: Check to use the restricted form of the mixed models (both fixed and random effects). The restricted model forces mixed interaction effects to sum to zero over the fixed effects. By default, Minitab fits the unrestricted model.

Balanced MANOVA – Graphs

Stat > ANOVA > Balanced MANOVA > Graphs

Displays residual plots. You do not have to store the residuals and fits in order to produce these plots.

Dialog box items

Residual Plots

Individual plots: Choose to display one or more plots.

Histogram of residuals: Check to display a histogram of the residuals.

Normal plot or residuals: Check to display a normal probability plot of the residuals.

Residuals versus fits: Check to plot the residuals versus the fitted values.

Residuals versus order: Check to plot the residuals versus the order of the data. The row number for each data point is shown on the x-axis--for example, 1 2 3 4... n.

Four in one: Choose to display a layout of a histogram of residuals, a normal plot of residuals, a plot of residuals versus fits, and a plot of residuals versus order.

Residuals versus the variables: Enter one or more columns containing the variables against which you want to plot the residuals. Minitab displays a separate graph for each column.

Balanced MANOVA – Results

Stat > ANOVA > Balanced MANOVA > Results

You can control the Session window output.

Dialog box items

Display of Results

Matrices (hypothesis, error, partial correlations): Check to display the hypothesis matrix H, the error matrix E, and a matrix of partial correlations. See MANOVA tests.

Eigen analysis: Check to display the eigenvalues and eigenvalues for the matrix E**-1 H.

Univariate analysis of variance: Check to perform a univariate analysis of variance for each response variable.

Expected mean squares for univariate analysis: Check to display the expected mean squares when you have requested univariate analysis of variance.

Display means corresponding to the terms: Display a table of means corresponding to specified terms from the model. For example, if you specify A B D A ∗ B ∗ D, four table of means will be printed, one for each main effect, A, B, D, and one for the three-way interaction, A ∗ B ∗ D.

Custom multivariate tests for the following terms: Perform 4 multivariate tests for model terms that you specify. See Specifying terms to test. Default tests are performed for all model terms.

Error: Designate an error term for the four multivariate tests. It must be a single term that is in the model. If you do not specify an error term, Minitab uses the error associated with mean squares error, as in the univariate case.

Specifying terms to test – Balanced MANOVA

In the **Results** subdialog box, you can specify model terms in **Custom multivariate test for the following terms** and designate an error term in **Error** and Minitab will perform four multivariate tests for those terms. This option is probably less useful for balanced MANOVA than it is for general MANOVA; because you can specify factors to be random with balanced MANOVA, Minitab will use the correct error terms. This option exists for special purpose tests.

If you specify an error term, it must be a single term that is in the model. This error term is used for all requested tests. If you do not specify an error term, Minitab determines an appropriate error term.

MANOVA tests – Balanced MANOVA

Minitab automatically performs four multivariate tests–Wilks' test, Lawley-Hotelling test, Pillai's test, and Roy's largest root test–for each term in the model and for specially requested terms (see Specifying terms to test). All four tests are based on two SSCP (sums of squares and cross products) matrices: H, the hypothesis matrix and E, the error matrix. There is one H associated with each term. E is the matrix associated with the error for the test. These matrices are displayed when you request the hypothesis matrices and are labeled by SSCP Matrix.

The test statistics can be expressed in terms of either H and/or E or the eigenvalues of E**-1 H. You can request to have these eigenvalues printed. (If the eigenvalues are repeated, corresponding eigenvectors are not unique and in this case, the eigenvectors Minitab prints and those in books or other software may not agree. The MANOVA tests, however, are always unique.)

You can also display the matrix of partial correlations, which are the correlations among the residuals, or alternatively, the correlations among the responses conditioned on the model. The formula for this matrix is W**-.5 E W**-.5, where E is the error matrix and W has the diagonal of E as its diagonal and 0's off the diagonal.

Hotelling's T-Squared Test

Hotelling's T-squared test to compare the mean vectors of two groups is a special case of MANOVA, using one factor that has two levels. Minitab's MANOVA option can be used to do this test. The usual T-squared test statistic can be calculated from Minitab's output using the relationship T-squared = (N-2) U, where N is the total number of observations and U is the Lawley-Hotelling trace. S, the pooled covariance matrix, is E / (N-2), where E is the error matrix.

Balanced MANOVA – Storage

Stat > ANOVA > Balanced MANOVA > Storage

Store fits and residuals for each response. If you fit a full model, fits are cell means. If you fit a reduced model, fits are least squares estimates.

Dialog box items

Fits: Check to store the fitted values for each observation in the data set in the next available columns, using one column for each response.

Residuals: Check to store the residuals using one column for each response.

Testing the equality of means from multiple responses

Balanced MANOVA and general MANOVA are procedures for testing the equality of vectors of means from multiple responses. Your choice between these two procedures depends upon the experimental design and the available options. Both procedures can fit MANOVA models to balanced data with up to 31 factors.

- Balanced MANOVA is used to perform multivariate analysis of variance with balanced designs. See Balanced designs. You can also specify factors to be random and obtain expected means squares. Use general MANOVA with unbalanced designs.
- General MANOVA is used to perform multivariate analysis of variance with either balanced or unbalanced designs that can also include covariates. You cannot specify factors to be random as you can for balanced MANOVA, although you can work around this restriction by specifying the error term for testing different model terms.

The table below summarizes the differences between Balanced and General MANOVA:

	Balanced MANOVA	General MANOVA
Can fit unbalanced data	no	yes
Can specify factors as random and obtain expected means squares	yes	no
Can fit covariates	no	yes
Can fit restricted and unrestricted forms of a mixed model	yes	no; unrestricted only

Example of Balanced MANOVA

You perform a study in order to determine optimum conditions for extruding plastic film. You measure three responses–tear resistance, gloss, and opacity–five times at each combination of two factors–rate of extrusion and amount of an additive–each set at low and high levels. The data and example are from Johnson and Wichern [5], page 266. You use Balanced MANOVA to test the equality of means because the design is balanced.

1. Open the file EXH_MVAR.MTW.
2. Choose **Stat > ANOVA > Balanced MANOVA**.
3. In **Responses**, enter *Tear Gloss Opacity*.
4. In **Model**, enter *Extrusion* | *Additive*.
5. Click **Results**. Under **Display of Results**, check **Matrices (hypothesis, error, partial correlations)** and **Eigen analysis**.
6. Click **OK** in each dialog box.

Session window output

ANOVA: Tear, Gloss, Opacity versus Extrusion, Additive

```
MANOVA for Extrusion
s = 1    m = 0.5    n = 6.0

                       Test             DF
Criterion         Statistic      F  Num  Denom      P
Wilks'              0.38186  7.554    3     14  0.003
Lawley-Hotelling    1.61877  7.554    3     14  0.003
Pillai's            0.61814  7.554    3     14  0.003
Roy's               1.61877


SSCP Matrix for Extrusion

            Tear   Gloss  Opacity
Tear       1.740  -1.505   0.8555
Gloss     -1.505   1.301  -0.7395
```

```
Opacity   0.855  -0.739   0.4205


SSCP Matrix for Error

           Tear    Gloss  Opacity
Tear      1.764   0.0200   -3.070
Gloss     0.020   2.6280   -0.552
Opacity  -3.070  -0.5520   64.924


Partial Correlations for the Error SSCP Matrix

             Tear     Gloss   Opacity
Tear      1.00000   0.00929  -0.28687
Gloss     0.00929   1.00000  -0.04226
Opacity  -0.28687  -0.04226   1.00000


EIGEN Analysis for Extrusion


Eigenvalue  1.619  0.00000  0.00000
Proportion  1.000  0.00000  0.00000
Cumulative  1.000  1.00000  1.00000


Eigenvector        1       2        3
Tear          0.6541  0.4315   0.0604
Gloss        -0.3385  0.5163   0.0012

Opacity       0.0359  0.0302  -0.1209


MANOVA for Additive
s = 1    m = 0.5    n = 6.0

                       Test             DF
Criterion         Statistic      F  Num  Denom       P
Wilks'              0.52303  4.256    3     14  0.025
Lawley-Hotelling    0.91192  4.256    3     14  0.025
Pillai's            0.47697  4.256    3     14  0.025
Roy's               0.91192


SSCP Matrix for Additive

           Tear   Gloss  Opacity
Tear     0.7605  0.6825    1.931
Gloss    0.6825  0.6125    1.732
Opacity  1.9305  1.7325    4.901


EIGEN Analysis for Additive


Eigenvalue  0.9119  0.00000  0.00000
Proportion  1.0000  0.00000  0.00000
Cumulative  1.0000  1.00000  1.00000


Eigenvector        1        2        3
Tear         -0.6330   0.4480  -0.1276
Gloss        -0.3214  -0.4992  -0.1694
Opacity      -0.0684   0.0000   0.1102


MANOVA for Extrusion*Additive
s = 1    m = 0.5    n = 6.0

                       Test             DF
```

```
Criterion         Statistic      F  Num  Denom      P
Wilks'              0.77711  1.339    3     14  0.302
Lawley-Hotelling    0.28683  1.339    3     14  0.302
Pillai's            0.22289  1.339    3     14  0.302
Roy's               0.28683


SSCP Matrix for Extrusion*Additive

            Tear    Gloss  Opacity
Tear     0.000500  0.01650  0.04450
Gloss    0.016500  0.54450  1.46850
Opacity  0.044500  1.46850  3.96050


EIGEN Analysis for Extrusion*Additive


Eigenvalue  0.2868  0.00000  0.00000

Proportion  1.0000  0.00000  0.00000
Cumulative  1.0000  1.00000  1.00000


Eigenvector        1        2        3
Tear         -0.1364   0.1806   0.7527
Gloss        -0.5376  -0.3028  -0.0228
Opacity      -0.0683   0.1102  -0.0000
```

Interpreting the results

By default, Minitab displays a table of the four multivariate tests (Wilks', Lawley-Hotelling, Pillai's, and Roy's) for each term in the model. The values s, m, and n are used in the calculations of the F-statistics for Wilks', Lawley-Hotelling, and Pillai's tests. The F-statistic is exact if s = 1 or 2, otherwise it is approximate [6]. Because you requested the display of additional matrices (hypothesis, error, and partial correlations) and an eigen analysis, this information is also displayed. The output is shown only for one model term, Extrusion, and not for the terms Additive or Extrusion*Additive.

Examine the p-values for the Wilks', Lawley-Hotelling, and Pillai's test statistic to judge whether there is significant evidence for model effects. These values are 0.003 for the model term Extrusion, indicating that there is significant evidence for Extrusion main effects at a levels greater than 0.003. The corresponding p-values for Additive and for Additive*Extrusion are 0.025 and 0.302, respectively (not shown), indicating that there is no significant evidence for interaction, but there is significant evidence for Extrusion and Additive main effects at a levels of 0.05 or 0.10.

You can use the SSCP matrices to assess the partitioning of variability in a similar way as you would look at univariate sums of squares. The matrix labeled as SSCP Matrix for Extrusion is the hypothesis sums of squares and cross-products matrix, or H, for the three response with model term Extrusion. The diagonal elements of this matrix, 1.740, 1.301, and 0.4205, are the univariate ANOVA sums of squares for the model term Extrusion when the response variables are Tear, Gloss, and Opacity, respectfully. The off-diagonal elements of this matrix are the cross products.

The matrix labeled as SSCP Matrix for Error is the error sums of squares and cross-products matrix, or E. The diagonal elements of this matrix, 1.764, 2.6280, and 64.924, are the univariate ANOVA error sums of squares when the response variables are Tear, Gloss, and Opacity, respectfully. The off-diagonal elements of this matrix are the cross products. This matrix is displayed once, after the SSCP matrix for the first model term.

You can use the matrix of partial correlations, labeled as Partial Correlations for the Error SSCP Matrix, to assess how related the response variables are. These are the correlations among the residuals or, equivalently, the correlations among the responses conditioned on the model. Examine the off-diagonal elements. The partial correlations between Tear and Gloss of 0.00929 and between Gloss and Opacity of -0.04226 are small. The partial correlation of -0.28687 between Tear and Opacity is not large. Because the correlation structure is weak, you might be satisfied with performing univariate ANOVA for these three responses. This matrix is displayed once, after the SSCP matrix for error.

You can use the eigen analysis to assess how the response means differ among the levels of the different model terms. The eigen analysis is of E-1 H, where E is the error SCCP matrix and H is the response variable SCCP matrix. These are the eigenvalues that are used to calculate the four MANOVA tests.

Place the highest importance on the eigenvectors that correspond to high eigenvalues. In the example, the second and third eigenvalues are zero and therefore the corresponding eigenvectors are meaningless. For both factors, Extrusion and Additive, the first eigenvectors contain similar information The first eigenvector for Extrusion is 0.6541, -0.3385, 0.0359 and for Additive it is -0.6630, -0.3214, -0.0684 (not shown). The highest absolute value within these eigenvectors is for the response Tear, the second highest is for Gloss, and the value for Opacity is small. This implies that the Tear means have the largest differences between the two factor levels of either Extrusion or Additive, the Gloss means have the next largest differences, and the Opacity means have small differences.

General MANOVA

General MANOVA

Stat > ANOVA > General MANOVA

Use general MANOVA to perform multivariate analysis of variance (MANOVA) with balanced and unbalanced designs, or if you have covariates. This procedure takes advantage of the data covariance structure to simultaneously test the equality of means from different responses.

Calculations are done using a regression approach. A "full rank" design matrix is formed from the factors and covariates and each response variable is regressed on the columns of the design matrix.

Factors may be crossed or nested, but they cannot be declared as random; it is possible to work around this restriction by specifying the error term to test model terms (See Specifying terms to test). Covariates may be crossed with each other or with factors, or nested within factors. You can analyze up to 50 response variables with up to 31 factors and 50 covariates at one time.

Dialog box items

Responses: Enter up to 50 numeric columns containing the response variables.

Model: Type the model terms that you want to fit.

<Covariates>

<Options>

<Graphs>

<Results>

<Storage>

Data – General MANOVA

Set up your worksheet in the same manner as with balanced MANOVA: one column for each response variable, one column for each factor, and one column for each covariate, so that there is one row of the worksheet for each observation. The factor columns may be numeric, text, or date/time. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. (See Ordering Text Categories.) You may include up to 50 response variables and up to 31 factors at one time.

Although models can be unbalanced in general MANOVA, they must be "full rank." That is, there must be enough data to estimate all the terms in your model. For example, suppose you have a two-factor crossed model with one empty cell. Then you can fit the model with terms A B, but not A B A∗B. Minitab will tell you if your model is not full rank. In most cases, eliminating some of the high order interactions in your model (assuming, of course, they are not important) can solve non-full rank problems.

Nesting does not need to be balanced. If factor B is nested within factor A, there can be unequal levels of B within each level of A. In addition, the subscripts used to identify the B levels can differ within each level of A.

If any response, factor, or covariate column contains missing data, that entire observation (row) is excluded from all computations. If an observation is missing for one response variable, that row is eliminated for all responses.

To perform a general MANOVA

1 Choose **Stat > ANOVA > General MANOVA**.

2 In **Responses**, enter up to 50 numeric columns containing the response variables.

3 In **Model**, type the model terms that you want to fit. See Overview of Balanced ANOVA and GLM.

4 If you like, use any dialog box options, then click **OK**.

General MANOVA – Covariates

Stat > ANOVA > General MANOVA > Covariates

Enter covariates into the model.

Dialog box items

Covariates: Enter up to 50 columns containing the covariates.

General MANOVA – Options

Stat > ANOVA > General MANOVA > Options

Allows you to perform a weighted regression.

Dialog box items

Do a weighted fit, using weights in: Enter a column containing weights to perform weighted regression.

General MANOVA – Graphs

Stat > ANOVA > General MANOVA > Graphs

Displays residual plots. You do not have to store the residuals and fits in order to produce these plots.

Dialog box items

Residuals for Plots: You can specify the type of residual to display on the residual plots.

Regular: Choose to plot the regular or raw residuals.

Standardized: Choose to plot the standardized residuals.

Deleted: Choose to plot the Studentized deleted residuals.

Residual Plots

Individual plots: Choose to display one or more plots.

Histogram of residuals: Check to display a histogram of the residuals.

Normal plot of residuals: Check to display a normal probability plot of the residuals.

Residuals versus fits: Check to plot the residuals versus the fitted values.

Residuals versus order: Check to plot the residuals versus the order of the data. The row number for each data point is shown on the x-axis--for example, 1 2 3 4... n.

Four in one: Choose to display a layout of a histogram of residuals, a normal plot of residuals, a plot of residuals versus fits, and a plot of residuals versus order.

Residuals versus the variables: Enter one or more columns containing the variables against which you want to plot the residuals. Minitab displays a separate graph for each column.

General MANOVA – Results

Stat > ANOVA > General MANOVA > Results

Display certain MANOVA and ANOVA output, display means, and customize MANOVA tests.

Dialog box items

Display of Results

Matrices (hypothesis, error, partial correlations): Check to display the hypothesis matrix H, the error matrix E, and a matrix of partial correlations. See MANOVA tests.

Eigen analysis: Check to display the eigenvalues and eigenvalues for the matrix E**-1 H.

Univariate analysis of variance: Check to perform a univariate analysis of variance for each response variable.

Display least squares means corresponding to the terms: Enter terms for which to display a table of means. For example, if you specify A B D A∗B∗D, four table of means will be displayed, one for each main effect, A, B, D, and one for the three-way interaction, A∗B∗D.

Custom multivariate tests for the following terms: Enter terms for which to perform 4 multivariate tests. See Specifying terms to test. By default the tests are performed for all model terms.

Error: Enter an error term for the four multivariate tests. It must be a single term that is in the model. If you do not specify an error term, Minitab uses the error associated with mean squares error, as in the univariate case.

Specifying terms to test

In the Results subdialog box, you can specify model terms in **Custom multivariate test for the following terms** and designate the error term in **Error**. Minitab will perform four multivariate tests (see MANOVA tests) for those terms. This option is most useful when you have factors that you consider as random factors. Model terms that are random or that are interactions with random terms may need a different error term than general MANOVA supplies. You can determine the appropriate error term by entering one response variable with General Linear Model, choose to display the expected mean square, and determine which error term was used for each model terms.

If you specify an error term, it must be a single term that is in the model. This error term is used for all requested tests. If you have different error terms for certain model terms, enter each separately and exercise the general MANOVA dialog for each one. If you do not specify an error term, Minitab uses MSE.

MANOVA tests – General MANOVA

The MANOVA tests with general MANOVA are similar to those performed for balanced MANOVA. See MANOVA tests for Balanced Designs for details.

However, with general MANOVA, there are two SSCP matrices associated with each term in the model, the sequential SSCP matrix and the adjusted SSCP matrix. These matrices are analogous to the sequential SS and adjusted SS in univariate General Linear Model. In fact, the univariate SS's are along the diagonal of the corresponding SSCP matrix. If you do not specify an error term in **Error** when you enter terms in **Custom multivariate tests for the following terms**, then the adjusted SSCP matrix is used for H and the SSCP matrix associated with MSE is used for E. If you do specify an error term, the sequential SSCP matrices associated with H and E are used. Using sequential SSCP matrices guarantees that H and E are statistically independent.

General MANOVA – Storage

Stat > ANOVA > General MANOVA > Storage

Stores the residuals, fitted values, and many other diagnostics for further analysis (see Checking your model).

Dialog box items

Storage

Coefficients: Check to store the coefficients for a model that corresponds to the design matrix. (If M1 contains the design matrix and C1 the coefficients, then M1 times C1 gives the fitted values.)

Fits: Check to store the fitted values.

Residuals: Check to store the residuals.

Standardized residuals: Check to store the standardized residuals.

Deleted t residuals: Check to store Studentized residuals.

Hi [leverage]: Check to store leverages.

Cook's distance: Check to store Cook's distance.

DFITS: Check to store DFITS.

Design matrix: Check to store the design matrix corresponding to your model.

Testing the equality of means from multiple responses

Balanced MANOVA and general MANOVA are procedures for testing the equality of vectors of means from multiple responses. Your choice between these two procedures depends upon the experimental design and the available options. Both procedures can fit MANOVA models to balanced data with up to 31 factors.

- Balanced MANOVA is used to perform multivariate analysis of variance with balanced designs. See Balanced designs. You can also specify factors to be random and obtain expected means squares. Use general MANOVA with unbalanced designs.
- General MANOVA is used to perform multivariate analysis of variance with either balanced or unbalanced designs that can also include covariates. You cannot specify factors to be random as you can for balanced MANOVA, although you can work around this restriction by specifying the error term for testing different model terms.

The table below summarizes the differences between Balanced and General MANOVA:

	Balanced MANOVA	General MANOVA
Can fit unbalanced data	no	yes
Can specify factors as random and obtain expected means squares	yes	no
Can fit covariates	no	yes
Can fit restricted and unrestricted forms of a mixed model	yes	no; unrestricted only

Test for Equal Variances

Test for Equal Variances

Stat > ANOVA > Test for Equal Variances

Use variance test to perform hypothesis tests for equality or homogeneity of variance using Bartlett's and Levene's tests. An F Test replaces Bartlett's test when you have just two levels.

Many statistical procedures, including analysis of variance, assume that although different samples may come from populations with different means, they have the same variance. The effect of unequal variances upon inferences depends in part upon whether your model includes fixed or random effects, disparities in sample sizes, and the choice of multiple comparison procedure. The ANOVA F-test is only slightly affected by inequality of variance if the model contains fixed factors only and has equal or nearly equal sample sizes. F-tests involving random effects may be substantially affected, however [19]. Use the variance test procedure to test the validity of the equal variance assumption.

Dialog box items

Response: Enter the column containing the response variable.

Factors: Enter the columns containing the factors in the model.

Confidence level: Enter a value from 0 to 100 for the level of confidence desired for the confidence intervals displayed on the graph. The default level is 95. Minitab uses the Bonferroni method to calculate the simultaneous confidence intervals.

Title: Type the desired text in this box to replace the default title with your own custom title.

<Storage>

Data – Test for Equal Variances

Set up your worksheet with one column for the response variable and one column for each factor, so that there is one row for each observation. Your response data must be in one column. You may have up to 9 factors. Factor columns may be numeric, text, or date/time, and may contain any value. If there are many cells (factors and levels), the print in the output chart can get very small.

Rows where the response column contains missing data (∗) are automatically omitted from the calculations. When one or more factor columns contain missing data, Minitab displays the chart and Bartlett's test results. When you have missing data in a factor column, Minitab displays the Levene's test results only when two or more cells have multiple observations and one of those cells has three or more observations.

Data limitations include the following:

1. If none of the cells have multiple observations, nothing is calculated. In addition, there must be at least one nonzero standard deviation
2. The F-test for 2 levels requires both cells to have multiple observations.
3. Bartlett's test requires two or more cells to have multiple observations.
4. Levene's test requires two or more cells to have multiple observations, but one cell must have three or more.

Bartlett's versus Levene's tests

Minitab calculates and displays a test statistic and p-value for both Bartlett's test and Levene's test where the null hypothesis is of equal variances versus the alternative of not all variances being equal. If there are only two levels, an F-test is performed in place of Bartlett's test.

- Use Bartlett's test when the data come from normal distributions; Bartlett's test is not robust to departures from normality.
- Use Levene's test when the data come from continuous, but not necessarily normal, distributions. This method considers the distances of the observations from their sample median rather than their sample mean, makes the test more robust for smaller samples.

To perform a test for equal variances

1. Choose **Stat > ANOVA > Test for Equal Variances**.
2. In **Response**, enter the column containing the response.
3. In **Factors**, enter up to nine columns containing the factor levels.
4. If you like, use any dialog box options, then click **OK**.

Test for Equal Variances – Storage

Stat > ANOVA > Test for Equal Variances > Storage

Allows for the storage of the cell standard deviations, cell variances, and confidence limits for the standard deviations.

Dialog box items

Storage

Standard deviations: Check to store the standard deviation of each cell (each level of each factor).

Variances: Check to store the variance of each cell.

Upper confidence limits for sigmas: Check to store the upper confidence limits for the standard deviations.

Lower confidence limits for sigmas: Check to store the lower confidence limits for the standard deviations.

Example of Performing a Test for Equal Variance

You study conditions conducive to potato rot by injecting potatoes with bacteria that cause rotting and subjecting them to different temperature and oxygen regimes. Before performing analysis of variance, you check the equal variance assumption using the test for equal variances.

1 Open the worksheet EXH_AOV.MTW.

2 Choose **Stat > ANOVA > Test for Equal Variances**.

3 In **Response**, enter *Rot*.

4 In **Factors**, enter *Temp Oxygen*. Click **OK**.

Session window output

Test for Equal Variances: Rot versus Temp, Oxygen

```
95% Bonferroni confidence intervals for standard deviations

Temp  Oxygen  N    Lower    StDev    Upper
  10       2  3  2.26029  5.29150   81.890
  10       6  3  1.28146  3.00000   46.427
  10      10  3  2.80104  6.55744  101.481
  16       2  3  1.54013  3.60555   55.799
  16       6  3  1.50012  3.51188   54.349
  16      10  3  3.55677  8.32666  128.862


Bartlett's Test (normal distribution)
Test statistic = 2.71, p-value = 0.744


Levene's Test (any continuous distribution)
Test statistic = 0.37, p-value = 0.858


Test for Equal Variances: Rot versus Temp, Oxygen
```

Graph window output

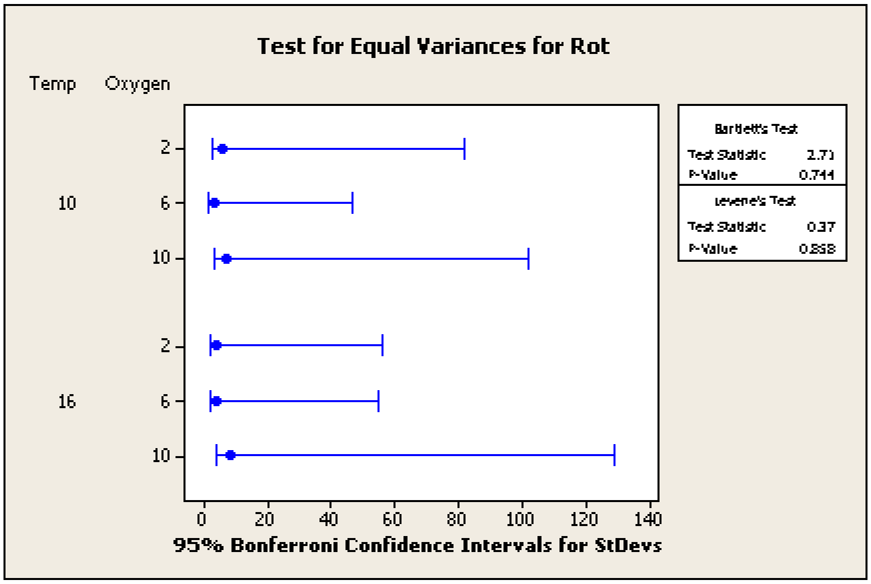

Interpreting the results

The test for equal variances generates a plot that displays Bonferroni 95% confidence intervals for the response standard deviation at each level. Bartlett's and Levene's test results are displayed in both the Session window and in the graph. Note that the 95% confidence level applies to the family of intervals and the asymmetry of the intervals is due to the skewness of the chi-square distribution.

For the potato rot example, the p-values of 0.744 and 0.858 are greater than reasonable choices of α, so you fail to reject the null hypothesis of the variances being equal. That is, these data do not provide enough evidence to claim that the populations have unequal variances.

Interval Plot

Interval Plot

Gallery

Data

One Y, Simple

To display a simple interval plot

Example, one y - simple

One Y, With Groups

To display an interval plot with groups

Example, one y - with groups

Multiple Y's, Simple

To display a simple interval plot with multiple y's

Example multiple y's - simple

Multiple Y's, With Groups

To display an interval plot with multiple y's and groups

Example, multiple y's - with groups

Main Effects Plot

Main Effects Plot

Stat > ANOVA > Main Effects Plot

Use Main Effects Plot to plot data means when you have multiple factors. The points in the plot are the means of the response variable at the various levels of each factor, with a reference line drawn at the grand mean of the response data. Use the main effects plot for comparing magnitudes of main effects.

Use Factorial Plots to generate main effects plots specifically for two-level factorial designs.

Dialog box items

Responses: Enter the columns containing the response data. You can include up to 50 responses.

Factors: Enter the columns containing the factor levels. You can include up to 9 factors.

<Options>

Data – Main Effects Plot

Set up your worksheet with one column for the response variable and one column for each factor, so that each row in the response and factor columns represents one observation. It is not required that your data be balanced.

The factor columns may be numeric, text, or date/time and may contain any values. If you wish to change the order in which text levels are processed, you can define your own order. See Ordering Text Categories. You may have up to 9 factors.

Missing values are automatically omitted from calculations.

To perform a main effects plot

1 Choose **Stat > ANOVA > Main Effects Plot**.

2 In **Responses**, enter the columns containing the response data.

3 In **Factors**, enter the columns containing the factor levels. You can enter up to 9 factors.

4 If you like, use any dialog box options, then click **OK**.

Main Effects Plot – Options

Stat > ANOVA > Main Effects Plot > Options

Allows you control the y scale minima and maxima and to add a title to the main effects plot.

Minimum for Y (response) scale: Enter either a single scale minimum for all responses or one scale minimum for each response.

Maximum for Y (response) scale: Enter either a single scale minimum for all responses or one scale minimum for each response.

Title: To replace the default title with your own custom title, type the desired text in this box.

Example of Main Effects Plot

You grow six varieties of alfalfa on plots within four different fields and you weigh the yield of the cuttings. You are interested in comparing yields from the different varieties and consider the fields to be blocks. You want to preview the data and examine yield by variety and field using the main effects plot.

1 Open the worksheet ALFALFA.MTW.

2 Choose **Stat > ANOVA > Main Effects Plot**.

3 In **Responses**, enter *Yield*.

4 In **Factors**, enter *Variety Field*. Click **OK**.

Graph window output

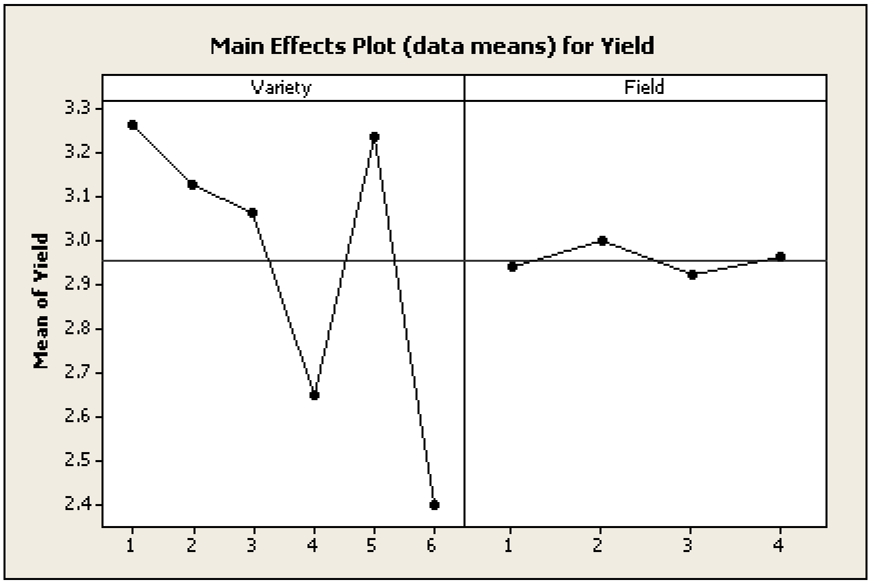

Interpreting the results

The main effects plot displays the response means for each factor level in sorted order if the factors are numeric or date/time or in alphabetical order if text, unless value ordering has been assigned (see Ordering Text Categories). A horizontal line is drawn at the grand mean. The effects are the differences between the means and the reference line. In the example, the variety effects upon yield are large compared to the effects of field (the blocking variable).

Interactions Plot

Interactions Plot

Stat > ANOVA > Interactions Plot

Interactions Plot creates a single interaction plot for two factors, or a matrix of interaction plots for three to nine factors. An interactions plot is a plot of means for each level of a factor with the level of a second factor held constant. Interactions plots are useful for judging the presence of interaction.

Interaction is present when the response at a factor level depends upon the level(s) of other factors. Parallel lines in an interactions plot indicate no interaction. The greater the departure of the lines from the parallel state, the higher the degree of interaction. To use interactions plot, data must be available from all combinations of levels.

Use Interactions plots for factorial designs to generate interaction plots specifically for 2-level factorial designs, such as those generated by Fractional Factorial Design, Central Composite Design, and Box-Behnken Design.

Dialog box items

Responses: Enter the columns containing the response data. You can include up to 50 responses.

Factors: Enter the columns containing the factor levels. You can include up to 9 factors.

Display full interaction plot matrix: Check to display the full interaction matrix when more than two factors are specified instead of displaying only the upper right portion of the matrix. In the full matrix, the transpose of each plot in the upper right displays in the lower left portion of the matrix. The full matrix takes longer to display than the half matrix.

<Options>

Data – Interactions Plot

Set up your worksheet with one column for the response variable and one column for each factor, so that each row in the response and factor columns represents one observation. Your data is not required to be balanced.

The factor columns may be numeric, text, or date/time and may contain any values. If you wish to change the order in which text levels are processed, you can define your own order. See Ordering Text Categories. You may have from 2 through 9 factors.

Missing data are automatically omitted from calculations.

To display an interactions plot

1 Choose **Stat > ANOVA > Interactions Plot**.

2 In **Responses**, enter the columns containing the response data.

3 In **Factors**, enter from 2 to 9 columns containing the factor levels. If you have two factors, the x-variable will be the second factor that you enter.

4 If you like, use any of the dialog box options, then click **OK**.

Interactions Plot – Options

Stat > ANOVA > Interactions Plot > Options

Allows you control the y scale minima and maxima and to add a title to the interaction plot.

Minimum for Y (response) scale: Enter either a single scale minimum for all responses or one scale minimum for each response.

Maximum for Y (response) scale: Enter either a single scale minimum for all responses or one scale minimum for each response.

Title: To replace the default title with your own custom title, type the desired text in this box.

Example of an Interactions Plot with Two Factors

You conduct an experiment to test the effect of temperature and glass type upon the light output of an oscilloscope (example and data from [13], page 252). There are three glass types and three temperatures, 100, 125, and 150 degrees Fahrenheit. You choose interactions plot to visually assess interaction in the data. You enter the quantitative variable second because you want this variable as the x variable in the plot.

1 Open the worksheet EXH_AOV.

2 Choose **Stat > ANOVA > Interactions Plot**.

3 In **Responses**, enter *LightOutput*.

4 In **Factors**, enter *GlassType Temperature*. Click **OK**.

Graph window output

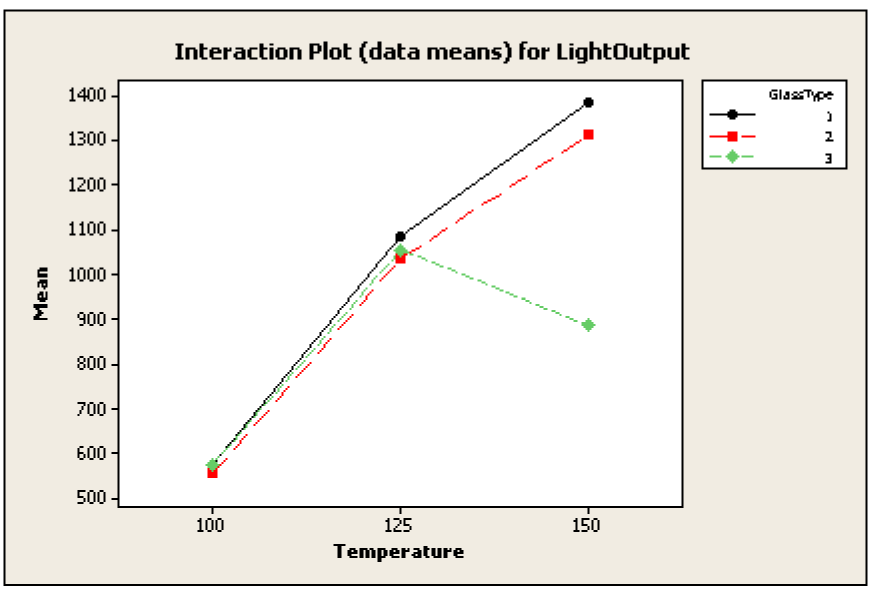

Interpreting the results

This interaction plot shows the mean light output versus the temperature for each of the three glass types. The legend shows which symbols and lines are assigned to the glass types. The means of the factor levels are plotted in sorted order if numeric or date/time or in alphabetical order if text, unless value ordering has been assigned (see Ordering Text Categories).

This plot shows apparent interaction because the lines are not parallel, implying that the effect of temperature upon light output depends upon the glass type. We test this using the General Linear Model.

Example of an Interactions Plot with more than Two Factors

Plywood is made by cutting thin layers of wood from logs as they are spun on their axis. Considerable force is required to turn a log hard enough so that a sharp blade can cut off a layer. Chucks are inserted into the ends of the log to apply the torque necessary to turn the log. You conduct an experiment to study factors that affect torque. These factors are diameter of the logs, penetration distance of the chuck into the log, and the temperature of the log. You wish to preview the data to check for the presence of interaction.

1 Open the worksheet PLYWOOD.MTW.

2 Choose **Stat > ANOVA > Interactions Plot**.

3 In **Responses**, enter *Torque*.

4 In **Factors**, enter *Diameter-Temp*. Click **OK**.

Graph window output

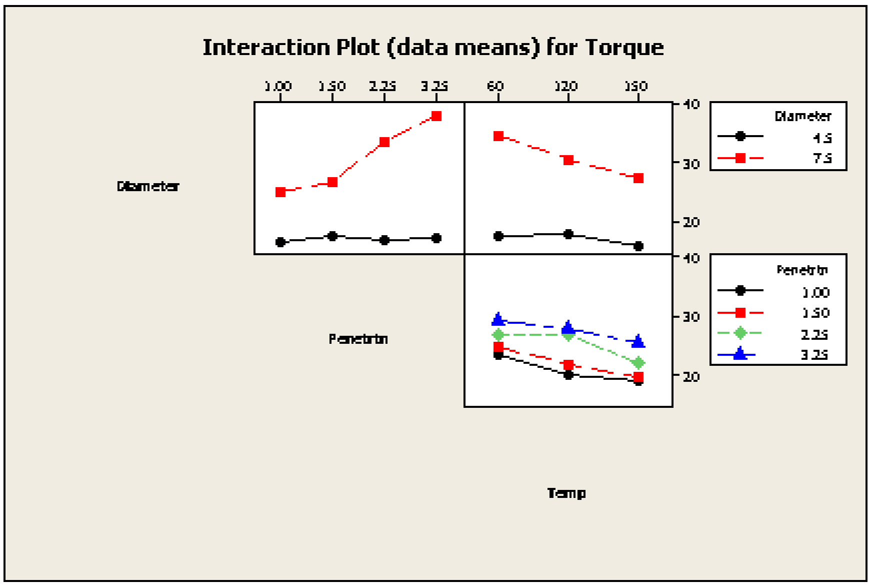

Interpreting the results

An interaction plot with three or more factors show separate two-way interaction plots for all two-factor combinations. In this example, the plot in the middle of the top row shows the mean torque versus the penetration levels for both levels of diameter, 4.5 and 7.5, averaged over all levels of temperature. There are analogous interactions plots for diameter by temperature (upper right) and penetration by temperature (second row).

For this example, the diameter by penetration and the diameter by temperature plots show nonparallel lines, indicating interaction. The presence of penetration by temperature interaction is not so easy to judge. This interaction might best be judged in conjunction with a model-fitting procedure, such as GLM.

Index